Palo Alto Networks Administrator's Guide

Release 5.0



Palo Alto Networks, Inc. www.paloaltonetworks.com © 2007-2015 Palo Alto Networks. All rights reserved. Palo Alto Networks, PAN-OS, and Panorama are trademarks of Palo Alto Networks, Inc. All other trademarks are the property of their respective owners. P/N 810-000107-00D

Table of Contents

Preface	13
About This Guide 1 Organization. 1 Typographical Conventions. 1 Notes and Cautions. 1 Related Documentation. 1	13 13 15 15
Chapter 1	
Introduction	17
Firewall Overview	17 18 19
Chapter 2 Getting Started	21
Preparing the Firewall	21 22 23 25 26 26 27 27 28 28 28 28
Chapter 3	~~
	29

System Setup, Configuration, and License Management
Defining Management Settings 30
Defining Operations Settings 37
Defining Services Settings 41

Defining Content ID Settings	43
Defining Session Settings	45
SNMP	47
Statistics Service	48
Comparing Configuration Files	49
Installing a License	50
Upgrading/Downgrading the PAN-OS Software	50
Upgrading PAN-OS in a High Availability Configuration	51
Downgrading PAN-OS Software	53
Maintenance Release Downgrade	53
Feature release Downgrade	54
Updating Threat and Application Definitions	55
Administrator Roles, Profiles, and Accounts.	56
Username and Password Requirements	57
Defining Administrator Roles	58
Defining Password Profiles	59
Creating Administrative Accounts	59
Specifying Access Domains for Administrators	61
Authentication Profiles	62
Setting Up Authentication Profiles	62
Creating a Local User Database	64
Configuring RADIUS Server Settings	65
Configuring LDAP Server Settings	66
Contiguring Kerberos Settings (Native Active Directory Authentication)	67
Authentication Sequence	67
Setting Up Authentication Sequences	68
Firewall Logs	68
Logging Configuration	70
Scheduling Log Exports	71
Defining Contiguration Log Settings	72
Defining System Log Settings	72
Defining HIP Match Log Settings	/3
	73
Canfiguring SNMD Trans Destingtions	74
	/5
Configuring Syslog Servers.	76
Custom Syslog Field Descriptions.	77
Contiguring Email Notification Settings.	83
Viewing Alarms	85
Configuring Netflow Settings	85
Importing, Exporting and Generating Security Certificates	86
Certificates	86
Default Trusted Certificate Authorities	88
Certificate Profile	89
OCSP Responder	90
Encrypting Private Keys and Passwords on the Firewall	90
Master Key and Diagnostic Settings	91
Updating Master Keys	91
High Availability	93
Active/Passive HA	93
Active / Active HA	93

Packet Flow
Deployment Options
NAT Considerations
Setting Up HA
Enabling HA on the Firewall 101
Virtual Systems
Communications Among Virtual Systems 111
Shared Gateways 112
Defining Virtual Systems 113
Configuring Shared Gateways 115
Defining Custom Response Pages 115
Viewing Support Information

Network C	onfiguration
Fir	rewall Deployment
	Virtual Wire Deployments
	laver 2 Deployments
	laver 3 Deployments
	Tap Mode Deployments
	Defining Virtual Wires
	Packet Content Modification 126
Fir	rewall Interfaces
	Viewing the Current Interfaces 128
	Configuring Laver 2 Interfaces
	Configuring Layer 2 Subinterfaces 129
	Configuring Layer 3 Interfaces.
	Configuring Layer 3 Subinterfaces
	Configuring Virtual Wire Interfaces
	Configuring Virtual Wire Subinterfaces
	Configuring Aggregate Interface Groups
	Configuring Aggregate Ethernet Interfaces
	Configuring VLAN Interfaces
	Configuring Loopback Interfaces
	Configuring Tunnel Interfaces
	Configuring Tap Interfaces
	Configuring HA Interfaces 150
Se	curity Zones
	Defining Security Zones 151
VL	AN Support
Vi	rtual Routers and Routing Protocols
• •	Politing Information Protocol 153
	Open Shortest Path First
	Border Gateway Protocol
	Multicast Routing
	Defining Virtual Routers 155
DH	ICP Server and Relay 171
Dr	NJ FIUXY

Network Profiles	174
Defining Interface Management Profiles	175
Defining Monitor Profiles	177
Defining Zone Protection Profiles	178

Policies and Security Profiles	183
Policies	183
Guidelines on Defining Policies	184
Specifying Users and Applications for Policies	186
Security Policies	187
Defining Security Policies	187
NAT Policies	190
Determining Zone Configuration in NAT and Security Policy	193
NAT Rule Ontions	193
Defining Network Address Translation Policies	194
NAT Policy Examples	105
NAT64	196
Policy-Rased Forwarding Policies	100
Decryption Policies	202
Application Override Policies	205
Custom Application Definition with Application Override	205
Defining Application Override Policies	205
Cantive Portal Policies	205
Defining Captive Portal Policies	207
DoS Protection Policies	200
Defining DoS Protection Policies	209
Security Profiles	211
Antivirus Profiles	212
Anti-snyware Profiles	213
Vulnerability Protection Profiles	215
IPI Filtering Profiles	217
File Blocking Profiles	220
Data Filtering Profiles	223
DoS Profiles	225
Other Policy Objects	226
Addresses and Address Groups	227
Defining Address Ranges	227
Defining Address Groups	230
Defining Regions	230
Applications and Application Groups	231
Defining Applications.	233
Custom Applications with Signatures	236
Defining Application Groups	238
Application Filters	238
Services	239
Service Groups	240
Data Patterns	240
Custom URL Categories	242
Dynamic Block Lists	242

Custom Spyware and Vulnerability Signatures
Defining Data Patterns
Defining Spyware and Vulnerability Signatures
Security Profile Groups
Log Forwarding
Decryption Profiles
Schedules

Reports and Logs 251
Using the Dashboard 252
Using the Application Command Center
Using App-Scope
Summary Report
Change Monitor Report
Threat Monitor Report
Threat Map Report
Network Monitor Report
Traffic Map Report
Viewing the Logs
Viewing Session Information
Working with Botnet Reports 267
Configuring the Botnet Report
Managing Botnet Reports
Managing PDF Summary Reports
Managing User Activity Reports
Managing Report Groups 272
Scheduling Reports for Email Delivery 273
Viewing Reports
Generating Custom Reports
Identifying Unknown Applications and Taking Action
Taking Action
Requesting an App-ID from Palo Alto Networks
Other Unknown Traffic
Taking Packet Captures 278

Chapter 7	
Configuring the Firewall for User	
Identification	281
Overview of User Identification 2	281
How User Identification Works	281
Identifying Users and Groups	282
How User-ID Components Interact	283
User-ID Agent	283
PAN-OS User Mapping	283
Terminal Services Agent	283
PAN-OS LDAP Group Query	284

User Identification Agents 284
Captive Portals
Configuring the Firewall for User Identification
PAN-OS User Mapping Configuration 291
Configuring PAN-OS User Mapping 292
Configure a Firewall to Share User Mapping Data
Setting Up the User-ID Agent 296
Installing the User-ID Agent
Configuring the User-ID Agent
Discovering Domain Controllers
Monitoring User-ID Agent Operation
Uninstalling and Upgrading the User-ID Agent
Setting Up the Terminal Services Agent 301
Installing or Upgrading the Terminal Server Agent on the Terminal Server . 302
Configuring the Terminal Server Agent on the Terminal Server
Uninstalling the Terminal Server Agent on the Terminal Server

Configuring IPSec Tunnels
Virtual Private Networks
VPN Tunnels
IPSec and IKE
IPSec and IKE Crypto Profiles
Setting Up IPSec VPNs
Defining IKE Gateways
Setting Up IPSec Tunnels
Defining IKE Crypto Profiles
Defining IPSec Crypto Profiles
Viewing IPSec Tunnel Status on the Firewall
Sample VPN Configuration
Existing Topology
New Topology
Configure the VPN Connection
VPN Connectivity Troubleshooting
GlobalProtect Large Scale VPN Deployment
Overview
Deploying a Large Scale VPN Network
Certificates and the OCSP Responder
Global Protect Gateway Configuration
GlobalProtect Portal Configuration
GlobalProtect Satellite Configuration
Dynamic Routing Protocols and Large Scale VPNs
Backing up a GlobalProtect Portal

GlobalProtect	;
erview	;
GlobalProtect Authentication	j
ting Up GlobalProtect	,
ting Up and Activating the GlobalProtect Agent	2
Setting Up the GlobalProtect Agent	;
	GlobalProtect 335 erview 335 GlobalProtect Authentication 336 ing Up GlobalProtect 337 ing Up and Activating the GlobalProtect Agent 357 Setting Up the GlobalProtect Agent 353

Chapter 10

Configuring Quality of Service	355
Firewall Support for QoS	355
Configuring QoS for Firewall Interfaces	356
Defining QoS Profiles	358
Defining QoS Policies	359
Displaying QoS Statistics	362

Chapter 11

Setting	Up a VM-Series Firewall 363
	Overview
	System Requirements and Limitations
	Requirements
	Limitations
	About Licensing the VM-Series Firewall
	Registering the VM-Series Firewall
	Installing and Licensing the VM-Series Firewall
	Troubleshooting

Setting	Up Panorama
	Overview
	Setting Up Panorama as a Virtual Appliance
	Installing Panorama
	Configuring the Panorama Network Interface
	Expanding the Log Storage Capacity
	Adding a Virtual Disk
	Setting Up Storage Partitions
	Setting up Panorama on an M-Series Appliance
	Performing Initial Setup
	Logging in to Panorama
	Changing the Default Password
	Configuring High Availability (HA)
	Switching the Logging Priority in an HA Pair

Central Device Ma	anagement Using
-------------------	-----------------

Panorama
Accessing the Panorama Web Interface
Using the Panorama Interface
Panorama Tab
Adding Devices
Defining Device Groups
Panorama Administrator Roles, Profiles, and Accounts
Defining Panorama Administrator Roles
Creating Panorama Administrative Accounts
Specifying Panorama Access Domains for Administrators
Device Groups
Working with Policies
Working with Objects 395
Working with Devices 397
Commit Operation in Panorama
Panorama Backward Compatibility 399
Templates
Configuring Panorama Templates
Adding a New Template
Configuring a Template
Overriding Template Settings
Logging
Concreting User Activity Penerts
Using Panorama for Log Collection 404
Deploving Distributed Log Collection
Managing Log Collectors
Defining Log Collector Groups
Viewing Firewall Deployment Information
Backing Up Firewall Configurations
Scheduling Configuration Exports
Upgrading the Panorama Software

onfiguring WildFire 421
About WildFire
Setting Up WildFire on the Firewall
Configuring WildFire Settings on the Firewall
Configuring WildFire Forwarding
WildFire Data Filtering Log 425
Using the WildFire Portal 425
Configuring Settings on the WildFire Portal
Viewing WildFire Reports

Appendix A

Custom Pages	
	Default Antivirus Response Page
	Default Application Block Page
	Default File Blocking Block Page
	Default URL Filtering Response Page
	Default Anti-spyware Download Response Page
	Default Decryption Opt-out Response Page
	Captive Portal Comfort Page434
	URL Filtering Continue and Override Page
	SSL VPN Login Page
	SSL Certificate Revoked Notify Page436

Appendix B

Application Categories	Subcategories,	Technologies, an	d Characteristics 437
------------------------	----------------	------------------	-----------------------

Application Categories and Subcategories	437
Application Technologies	439
Application Characteristics	439

Appendix C

Common Criteria/	Federal	Information	Processing	Standards	Support .	441

Enabling CC/FIPS Mode	41
CC/FIPS Security Functions	43

Appendix D

Open Source Licenses 445
Artistic License
BSD
GNU General Public License
GNU Lesser General Public License 452
MIT/X11
OpenSSH 458
PSF
PHP
Zlib
Index

Preface

This preface contains the following sections:

- "About This Guide" in the next section
- "Organization" on page 13
- "Typographical Conventions" on page 15
- "Notes and Cautions" on page 15
- "Related Documentation" on page 15

About This Guide

This guide describes how to administer the Palo Alto Networks firewall using the device's web interface.

This guide is intended for system administrators responsible for deploying, operating, and maintaining the firewall.

Organization

This guide is organized as follows:

- **Chapter 1, "Introduction"**—Provides an overview of the firewall.
- Chapter 2, "Getting Started"—Describes how to install the firewall.
- **Chapter 3, "Device Management"**—Describes how to perform basic system configuration and maintenance for the firewall, including how to configure a pair of firewalls for high availability, define user accounts, update the software, and manage configurations.
- **Chapter 4, "Network Configuration"**—Describes how to configure the firewall for your network, including routing configuration.
- **Chapter 5, "Policies and Security Profiles"**—Describes how to configure security policies and profiles by zone, users, source/destination address, and application.
- **Chapter 6, "Reports and Logs"**—Describes how to view the reports and logs provided with the firewall.

- **Chapter 7, "Configuring the Firewall for User Identification"**—Describes how to configure the firewall to identify the users who attempt to access the network.
- **Chapter 8, "Configuring IPSec Tunnels"**—Describes how to configure IP Security (IPSec) tunnels on the firewall.
- **Chapter 9, "Configuring GlobalProtect"**—Describes GlobalProtect, which allows secure login from client systems located anywhere in the world.
- **Chapter 10, "Configuring Quality of Service"**—Describes how to configure quality of service (QoS) on the firewall.
- **Chapter 12, "Setting Up Panorama"**—Describes how to install the centralized management system for the Palo Alto Networks firewall.
- **Chapter 13, "Central Device Management Using Panorama"**—Describes how to use Panorama to manage multiple firewalls.
- **Chapter 14, "Configuring WildFire"**—describes how to use WildFire for analysis and reporting on malware that traverses the firewall.
- **Appendix A, "Custom Pages"**—Provides HTML code for custom response pages to notify end users of policy violations or special access conditions.
- Appendix B, "Application Categories, Subcategories, Technologies, and Characteristics"—Contains a list of the application categories defined by Palo Alto Networks.
- Appendix C, "Common Criteria/Federal Information Processing Standards Support"— Describes firewall support for the Federal Information Processing Standards 140-2.
- Appendix D, "Open Source Licenses"—Includes information on applicable open source licenses.

Typographical Conventions

Convention	Meaning	Example
boldface	Names of commands, keywords, and selectable items in the web interface	Click Security to open the Security Rules page.
italics	Name of parameters, files, directories, or Uniform Resource Locators (URLs)	The address of the Palo Alto Networks home page is http://www.paloaltonetworks.com
courier font	Coding examples and text that you enter at the command prompt	Enter the following command: set deviceconfig system dns- settings
Click	Click the left mouse button	Click Administrators under the Devices tab.
Right-click	Click the right mouse button.	Right-click on the number of a rule you want to copy, and select Clone Rule .

This guide uses the following typographical conventions for special terms and instructions.

Notes and Cautions

This guide uses the following symbols for notes and cautions.

Symbol	Description
	NOTE Indicates helpful suggestions or supplementary information.
2	CAUTION Indicates actions that could cause loss of data.

Related Documentation

The following additional documentation is provided with the firewall:

- Getting Started Guide
- Quick Start
- Palo Alto Networks License Agreement and Warranty

You can find additional related documentation at https://live.paloaltonetworks.com/community/documentation.

Related Documentation

Chapter 1 Introduction

This chapter provides an overview of the firewall:

- "Firewall Overview" in the next section
- "Features and Benefits" on page 18
- "Management Interfaces" on page 19

Firewall Overview

The Palo Alto Networks firewall allows you to specify security policies based on accurate identification of each application seeking access to your network. Unlike traditional firewalls that identify applications only by protocol and port number, the firewall uses packet inspection and a library of application signatures to distinguish between applications that have the same protocol and port, and to identify potentially malicious applications that use non-standard ports.

For example, you can define security policies for specific applications, rather than rely on a single policy for all port 80 connections. For each identified application, you can specify a security policy to block or allow traffic based on the source and destination zones and addresses (IPv4 and IPv6). Each security policy can also specify security profiles to protect against viruses, spyware, and other threats.

Features and Benefits

The firewall provides granular control over the traffic allowed to access your network. The primary features and benefits include:

- **Application-based policy enforcement**—Access control by application is far more effective when application identification is based on more than just protocol and port number. High risk applications can be blocked, as well as high risk behavior, such as file-sharing. Traffic encrypted with the s Layer (SSL) protocol can be decrypted and inspected.
- User Identification (User-ID)—User-ID allows administrators to configure and enforce firewall policies based on users and user groups, instead of or in addition to network zones and addresses. The firewall can communicate with many directory servers, such as Microsoft Active Directory, eDirectory, SunOne, OpenLDAP, and most other LDAP based directory servers to provide user and group information to the firewall. This information can then be used to provide an invaluable method of providing secure application enablement that can be defined per user or group. For example, the administrator could allow one organization to use a web-based application, but no other organizations in the company would be able to use that application. You can also configure granular control of certain components of an application based on users and groups. Refer to "Configuring the Firewall for User Identification" on page 281.
- **Threat prevention**—Threat prevention services that protect the network from viruses, worms, spyware, and other malicious traffic can be varied by application and traffic source (refer to "Security Profiles" on page 211).
- URL filtering—Outbound connections can be filtered to prevent access to inappropriate web sites (refer to "URL Filtering Profiles" on page 217).
- **Traffic visibility**—Extensive reports, logs, and notification mechanisms provide detailed visibility into network application traffic and security events. The Application Command Center (ACC) in the web interface identifies the applications with the most traffic and the highest security risk (refer to "Reports and Logs" on page 251).
- Networking versatility and speed—The firewall can augment or replace your existing firewall, and can be installed transparently in any network or configured to support a switched or routed environment. Multi-gigabit speeds and a single-pass architecture provide all services with little or no impact on network latency.
- **GlobalProtect**—GlobalProtect provides security for client systems, such as laptops, that are used in the field by allowing easy and secure login from anywhere in the world.
- **Fail-safe operation**—High availability support provides automatic failover in the event of any hardware or software disruption (refer to "Enabling HA on the Firewall" on page 101).
- **Malware analysis and reporting**—WildFire provides detailed analysis and reporting on malware that traverses the firewall.
- VM-Series Firewall—Provides a virtual instance of PAN-OS positioned for use in a virtualized data center environment and particularly well suited for private and public cloud deployments. Installs on any x86 device that is capable of running VMware ESXi, without the need to deploy Palo Alto Networks hardware.

• **Management and Panorama**—Each firewall is managed through an intuitive web interface or a command-line interface (CLI), or all devices can be centrally managed through the Panorama centralized management system, which has a web interface very similar to the device web interface.

Management Interfaces

The firewall supports the following management interfaces. Refer to "Supported Browsers" on page 27 for a list of supported browsers.

- Web interface—Configuration and monitoring over HTTP or HTTPS from a web browser.
- **CLI**—Text-based configuration and monitoring over Telnet, Secure Shell (SSH), or the console port (refer to the *PAN-OS Command Line Interface Reference Guide*).
- **Panorama**—Palo Alto Networks product that provides web-based management, reporting, and logging for multiple firewalls. The Panorama interface is similar to the device web interface, with additional management functions included. Refer to "Setting Up Panorama" on page 371 for instructions on installing Panorama and "Central Device Management Using Panorama" on page 381 for information on using Panorama.
- Simple Network Management Protocol (SNMP)—Palo Alto Networks products support SNMPv2c and SNMPv3, read-only access over SNMP, and support for TRAPS. Refer to "Configuring SNMP Trap Destinations" on page 75).
- **Syslog**—Provides message generation for one or more remote syslog servers (refer to "Configuring Syslog Servers" on page 76).
- XML API—Provides a Representational State Transfer (REST)-based interface to access device configuration, operational status, reports, and packet captures from the firewall. There is an API browser available on the firewall at *https://<firewall>/api*, where *<firewall>* is the host name or IP address of the firewall. This link provides help on the parameters required for each type of API call. An XML API usage guide is available on the DevCenter online community at *http://live.paloaltonetworks.com*.

Management Interfaces

Chapter 2 Getting Started

This chapter describes how to set up and start using the firewall:

- "Preparing the Firewall" in the next section
- "Setting Up the Firewall" on page 22
- "Using the Firewall Web Interface" on page 23
- "Getting Help Configuring the Firewall" on page 28



Note: Refer to "Setting Up Panorama" on page 371 for instructions on installing the Panorama centralized management system.

Preparing the Firewall

Perform the following tasks to prepare the firewall for setup:

- 1. Mount the firewall in a rack and power it up as described in the *Hardware Reference Guide*.
- 2. Register your firewall at *https://support.paloaltonetworks.com* to obtain the latest software and App-ID updates, and to activate support or subscriptions with the authorization codes emailed to you.
- 3. Obtain an IP address from your network administrator for configuring the management port on the firewall.

Setting Up the Firewall

To perform the initial firewall setup:

- 1. Connect your computer to the management port (MGT) on the firewall using an RJ-45 Ethernet cable.
- 2. Start your computer. Assign a static IP address to your computer on the 192.168.1.0 network (for example, 192.168.1.5) with a netmask of 255.255.255.0.
- 3. Launch a supported web browser and enter https://192.168.1.1.

The browser automatically opens the Palo Alto Networks login page.

- 4. Enter **admin** in both the **Name** and **Password** fields, and click **Login**. The system presents a warning that the default password should be changed. Click **OK** to continue.
- 5. On the **Device** tab, choose **Setup** and configure the following (for general instructions on configuring settings in the web interface, refer to "Using the Firewall Web Interface" on page 23):
 - On the **Management** tab under **Management Interface Settings**, enter the firewall's IP address, netmask, and default gateway.
 - On the Services tab, enter the IP address of the Domain Name System (DNS) server. Enter the IP address or host and domain name of the Network Time Protocol (NTP) server and select your time zone.
 - Click Support on the side menu.
 If this is the first Palo Alto Networks firewall for your company, click Register Device to register the firewall. (If you have already registered a firewall, you have received a user name and password.)
 Click the Activate support using authorization codes link and enter the authorization codes that have been emailed to you for any optional features. Use a space to separate multiple authorization codes.
- 6. Click Administrators under the Devices tab.
- 7. Click admin.
- 8. In the **New Password** and **Confirm New Password** fields, enter and confirm a casesensitive password (up to 15 characters).
- 9. Click **OK** to submit the new password.
- 10. Commit the configuration to make these settings active. When the changes are committed, the firewall will be reachable through the IP address assigned in Step 5. For information on committing changes, refer to "Committing Changes" on page 25.



Note: The default configuration of the firewall when delivered from the factory, or after a factory reset is performed, is a virtual wire between Ethernet ports 1 and 2 with a default policy to deny all inbound traffic and allow all outbound traffic.

Using the Firewall Web Interface

The following conventions apply when using the firewall interface.

• To display the menu items for a general functional category, click the tab, such as **Objects** or **Device**, near the top of the browser window.

Dashboard	ACC	Monitor	Policies	Objects	Network	1	Device	١
						_		2

• Click an item on the side menu to display a panel.



• To display submenu items, click the 主 icon to the left of an item. To hide submenu items, click the 📄 icon to the left of the item.



• On most configuration pages, you can click **Add** to create a new item.



• To delete one or more items, select their check boxes and click **Delete**. In most cases, the system prompts you to confirm by clicking **OK** or to cancel the deletion by clicking **Cancel**.

+ Add - Delete

• On some configuration pages, you can select the check box for an item and click **Clone** to create a new item with the same information as the selected item.

+ Add Delete O Clone

• To modify an item, click its underlined link.

	Name	Location	Protocol
	service-http	Predefined	TCP
¢	service-https	Predefined	TCP

• To view help information on a page, click the **Help** icon in upper right area of the page.



• To view the current list of tasks, click the **Tasks** icon in the lower right corner of the page. The Task Manager window opens to show the list of tasks, along with status, start times, associated messages, and actions. Use the **Show** drop-down list to filter the list of tasks.

Туре	Status	Start Time	Messages
Commit	Completed	09/09/1107:57:23	 In virtual-router vr.1: address 10.40.1.1/24 on interface ethernet1/7 is duplicate with address 10.40.1.2/24 on interface ethernet1/8. (Module: routed) Commit failed
Commit	Completed	09/09/11 07:56:16	 In virtual-router vr1: address 10.40.1.1/24 on interface ethernet1/7 is duplicate with address 10.400.1.2/24 on interface ethernet1/8. (Module: routed) Commit failed
Commit	Completed	09/09/11 07:52:26	
Commit	Completed	09/09/11 07:49:25	 Interface ethernet1/7 has no ip pool(Module: dhcpd) Configuration committed successfully
Auto Commit	Completed	09/08/11 15:06:15	 Configuration committed successfully Successfully committed last configuration L3 Service Configuration is changed. I3svc will be restarted. (Module: device)
<			

• The web interface language is controlled by the current language of the computer that is managing the device if a specific language preference has not been defined. For example, if the computer you use to manage the firewall has a locale of Spanish, when you log in to the firewall, the web interface will be in Spanish.

To specify a language that will always be used for a given account regardless of the locale of the computer, click the **Language** icon in the lower right corner of the page and the Language Preference window opens. Click the drop-down list to select the desired language and then click **OK** to save your change.

Language Preference 📀		
Language		
	es (Spanish)	٦
	en (English)	
	fr (French)	
	zh_TW (Chinese-Traditional)	
	ja (Japanese)	
	zh_CN (Chinese-Simplified)	

• On pages that list information you can modify (for example, the **Setup** page on the **Devices** tab), click the icon in the upper right corner of a section to edit the settings.



• After you configure settings, you must click **OK** or **Save** to store the changes. When you click **OK**, the current "candidate" configuration is updated.

Committing Changes

Click **Commit** at the top of the web interface to open the commit dialog box.

Commit 💿	📥 Commit
Doing a commit will overwrite the running configuration. Do you want to continue?	
Advanced	
Include Device and Network configuration	
✓ Include Shared Object configuration	
☑ Include Virtual System configuration	
All virtual systems	
Select one or more virtual systems	
PM VSYS	
vsys2	
Preview Changes OK Cancel	

The following options are available in the commit dialog box. Click the **Advanced** link, if needed, to display the options:

- **Include Device and Network configuration**—Include the device and network configuration changes in the commit operation.
- **Include Shared Object configuration**—(Multi-virtual system firewalls only) Include the shared object configuration changes in the commit operation.
- **Include Policy and Objects**—(Non-multi-virtual system firewalls only) Include the policy and object configuration changes in the commit operation.
- Include virtual system configuration—Include all virtual systems or choose Select one or more virtual systems.

For more information about committing changes, refer to "Defining Operations Settings" on page 37.

 Preview Changes—Click this button to bring up a two-pane window that shows proposed changes in the candidate configuration compared to the current running configuration. You can choose the number of lines of context to display, or show all lines. Changes are color coded based on items that have been added, modified, or deleted.

The **Device > Config Audit** feature performs the same function, refer to "Comparing Configuration Files" on page 49.



Note: Configuration changes that span multiple configuration areas may require a full commit. For example, if you click **Commit** and only select the **Include Device and Network configuration** option, some items that you changed in the Device tab will not commit. This includes certificates and User-ID options as well as Server Profiles used for User-ID, such as an LDAP server profile. This can also occur if you perform a partial commit after importing a configuration. To commit these types of changes, do a full commit and select both **Include Device and Network configuration and Include Policy and Object configuration**.

Navigating to Configuration Pages

Each configuration section in this guide shows the menu path to the configuration page. For example, to reach the **Vulnerability Protection** page, choose the **Objects** tab and then choose **Vulnerability Protection** under **Security Profiles** in the side menu. This is indicated in this guide by the following path:

Objects > Security Profiles > Vulnerability Protection

Using Tables on Configuration Pages

The tables on configuration pages include sorting and column chooser options. Click a column header to sort on that column, and click again to change the sort order. Click the arrow to the right of any column and select check boxes to choose the columns to display.



Required Fields

Required fields are shown with a light yellow background. A message indicating that the field is required appears when you hover over or click in the field entry area.

Name	I
escription	This field is required

Locking Transactions

The web interface provides support for multiple administrators by allowing an administrator to lock a current set of transactions, thereby preventing configuration changes or commit operations by another administrator until the lock is removed. The following types of locks are supported:

- **Config lock**—Blocks other administrators from making changes to the configuration. This type of lock can be set globally or for a virtual system. It can be removed only by the administrator who set it or by a superuser on the system.
- **Commit Lock**—Blocks other administrators from committing changes until all of the locks have been released. This type of lock prevents collisions that can occur when two administrators are making changes at the same time and the first administrator finishes and commits changes before the second administrator has finished. The lock is released when the current changes are committed by the administrator who applied the lock, or it can be released manually.

Any administrator can open the lock window to view the current transactions that are locked, along with a timestamp for each.

To lock a transaction, click the unlocked icon a on the top bar to open the Locks dialog box. Click **Take a Lock**, select the scope of the lock from the drop-down list, and click **OK**. Add additional locks as needed, and then click **Close** to close the Lock dialog box.

The transaction is locked, and the icon on the top bar changes to a locked icon that shows the number of locked items in parentheses.



To unlock a transaction, click the locked icon B on the top bar to open the Locks window. Click the B icon for the lock that you want to remove, and click **Yes** to confirm. Click **Close** to close the Lock dialog box.

You can arrange to automatically acquire a commit lock by selecting the **Automatically acquire commit lock** check box in the Management area of the **Device Setup** page. Refer to "System Setup, Configuration, and License Management" on page 30.

Supported Browsers

The following web browsers are supported for access to the firewall web interface:

- Internet Explorer 7+
- Firefox 3.6+
- Safari 5+
- Chrome 11+

Getting Help Configuring the Firewall

Use the information in this section to obtain help on using the firewall.

Obtaining More Information

To obtain more information about the firewall, refer to the following:

- **General information**—Go to *http://www.paloaltonetworks.com*.
- **Online help**—Click **Help** in the upper-right corner of the web interface to access the online help system.
- **Collaborative area for customer/partner interaction to share tips, scripts, and signatures**—Go to *https://live.paloaltonetworks.com/community/devcenter.*

Technical Support

For technical support, use the following methods:

- Go to the KnowledgePoint online support community at *http://live.paloaltonetworks.com*.
- Go to *https://support.paloaltonetworks.com*.

Chapter 3 Device Management

This chapter describes how to perform basic system configuration and maintenance for the firewall and includes overviews of the virtual systems, high availability, and logging functions:

- "System Setup, Configuration, and License Management" in the next section
- "Comparing Configuration Files" on page 49
- "Installing a License" on page 50
- "Upgrading/Downgrading the PAN-OS Software" on page 50
- "Updating Threat and Application Definitions" on page 55
- "Administrator Roles, Profiles, and Accounts" on page 56
- "Authentication Profiles" on page 62
- "Authentication Sequence" on page 67
- "Certificate Profile" on page 89
- "Firewall Logs" on page 68
- "Configuring SNMP Trap Destinations" on page 75
- "Configuring Syslog Servers" on page 76
- "Configuring Email Notification Settings" on page 83
- "Viewing Alarms" on page 85
- "Configuring Netflow Settings" on page 85
- "Importing, Exporting and Generating Security Certificates" on page 86
- "High Availability" on page 93
- "Virtual Systems" on page 110
- "Defining Custom Response Pages" on page 115
- "Viewing Support Information" on page 117

System Setup, Configuration, and License Management

The following sections describe how to define the network settings and manage configurations for the firewall:

- "Defining Management Settings" in the next section
- "Defining Operations Settings" on page 37
- "Defining Services Settings" on page 41
- "Defining Content ID Settings" on page 43
- "Defining Session Settings" on page 45
- "SNMP" on page 47
- "Statistics Service" on page 48
- "Installing a License" on page 50



Note: Refer to "Configuring WildFire" on page 421 for information on configuring the settings on the WildFire tab.

Defining Management Settings

Device > Setup > Management

The **Setup** page allows you to configure the firewall for management, operations, services, content identification, WildFire malware analysis and reporting, and session behavior.

If you do not want to use the management port, you can define a loopback interface and manage the firewall through the IP address of the loopback interface (refer to "Configuring Loopback Interfaces" on page 146).

Perform any of the following operations on this page:

• To change the host name or network settings, click **Edit** on the first table on the page, and specify the following information.

ltem	Description
General Settings	
Hostname	Enter a host name (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Domain	Enter the Fully Qualified Domain Name (FQDN) of the firewall (up to 31 characters).
Login Banner	Enter custom text that will be displayed on the firewall login page. The text is displayed below the Name and Password fields.

Table 1. Management Settings

Item	Description
Time Zone	Select the time zone of the firewall.
Locale	Select a language for PDF reports from the drop-down list. Refer to "Managing PDF Summary Reports" on page 270.
	If you have a specific language preference set for the web interface, PDF reports will still use the language specified in this locale setting. Refer to language preference in "Using the Firewall Web Interface" on page 23.
Time	To set the date and time on the firewall, click Set Time . Enter the current date in (YYYY/MM/DD) or click the calendar icon 🔄 to select a month and day. Enter the current time in 24-hour format (HH:MM:SS). You can also define an NTP server from Device > Setup > Services.
Serial Number	(Panorama only) Enter the serial number of the firewall.
Geo Location	Enter the latitude (-90.0 to 90.0) and longitude (-180.0 to 180.0) of the firewall.
Automatically acquire commit lock	Automatically apply a commit lock when you change the candidate configuration. For more information, refer to "Locking Transactions" on page 27.
Certificate Expiration Check	Instruct the firewall to create warning messages when on-box certificates near their expiration dates.
Multi Virtual System Capability	To enable the use of multiple virtual systems (if supported on the firewall model), click Edit for Multi Virtual System Capability near the top of the Setup page. Select the check box, and click OK . For more information about virtual systems, refer to "Virtual Systems" on page 110.
Authentication Settings	
Authentication Profile	Select the authentication profile to use for administrator access to the firewall. For instructions on configuring authentication profiles, refer to "Setting Up Authentication Profiles" on page 62.
Certificate Profile	Select the certificate profile to use for administrator access to the firewall. For instructions on configuring certificate profiles, refer to "Certificate Profile" on page 89.
Idle Timeout	Enter the timeout interval (1 - 1440 minutes). A value of 0 means that the management, web, or CLI session does not time out.
# Failed Attempts	Enter the number of failed login attempts that are allowed for the web interface and CLI before the account is locked. (1-10, default 0). 0 means that there is no limit.
Lockout Time	Enter the number of minutes that a user is locked out (0-60 minutes) if the number of failed attempts is reached. The default 0 means that there is no limit to the number of attempts.

Table 1. Management Settings (Continued)

ltem	Description
Panorama Settings	
Panorama Servers	Enter the IP address of Panorama, the Palo Alto Networks centralized management system (if any). The server address is required to manage the device using Panorama. If Panorama is in an HA configuration, enter the secondary Panorama server IP address in the second Panorama Servers field.
	Note: To remove any policies that Panorama propagates to managed firewalls, click the Disable Panorama Policy and Objects link. To keep a local copy of the policies and objects to your device before removing them from Panorama, click the Import Panorama Policy and Objects before disabling check box in the dialog box that opens. Click OK .
	Note: When you select the import check box, the policies and objects will be copied to the current candidate configuration. If you commit this configuration, the policies and objects will become part of your configuration and will no longer be managed by Panorama.
	To remove device and network templates, click the Disable Device and Network Template link. To keep a local copy of the device and network templates, click the Import Device and Network Templates before disabling check box in the dialog box that opens and click OK . When you select the import check box, the configuration defined in the device and network templates will be copied to the current candidate configuration. If you commit that configuration, these items will become part of your configuration and will no longer be managed by Panorama. Templates will no longer be accepted on the device until you click Enable Device and Network Templates .
Receive Timeout for connection to device/	 Device—Enter the timeout for receiving TCP messages from all managed devices (1-240 seconds, default 240). Panorama Enter the timeout for receiving TCP messages from Pan
Panorama	orama (1-240 seconds, default 240).
Send Timeout for	• Device—Enter the timeout for receiving TCP messages from all managed devices (1-240 seconds, default 240).
Panorama	• Panorama—Enter the timeout for receiving TCP messages from Pan- orama (1-240 seconds, default 240).
Retry Count for SSL	• Device—Enter the number of retries for attempts to send Secure Socket Layer (SSL) messages to managed devices (1-64, default 25).
Panorama	• Panorama—Enter the number of retries for attempts to send Secure Socket Layer (SSL) messages to Panorama (1-64, default 25).
Share Unused Address and Service Objects with Devices (Panorama only)	Select this check box to share all Panorama shared objects and device group specific objects with managed devices. When unchecked, Panorama policies are checked for references to address, address group, service, and service group objects and any objects that are not referenced will not be shared. This option will ensure that only necessary objects are being sent to managed devices in order to reduce the total object count.

Table 1. Management Settings (Continued)

ltem	Description
Shared Objects Take Precedence (Panorama only)	Select the check box to specify that shared objects take precedence over device group objects. This option is a system-wide setting and is off by default. When this option is off, device groups override corresponding objects of the same name. If the option is selected, device group objects cannot override corresponding objects of the same name from a shared location and any device group object with the same name as a shared object will be discarded.
Management Interface Settings	
MGT Interface Speed	Configure a data rate and duplex option for the management interface. The choices include 10Mbps, 100Mbps, and 1Gbps at full or half duplex. Use the default auto-negotiate setting to have the firewall determine the interface speed.
	This setting should match the port settings on the neighboring network equipment.
MGT Interface IP Address	Enter the IP address of the management port. Alternatively, you can use the IP address of a loopback interface for device management. This address is used as the source address for remote logging.
Netmask	Enter the network mask for the IP address, such as "255.255.255.0".
Default Gateway	Enter the IP address of the default router (must be on the same subnet as the management port).
MGT Interface IPv6 Address	(Optional) Enter the IPv6 address of the management port. An IPv6 prefix length is required to indicate the netmask, for example 2001:400:f00::1/64.
Default IPv6 Gateway	Enter the IPv6 address of the default router (must be on the same subnet as the management port), if you assigned an IPv6 address to the management port.
MGT Interface Services	Select the services enabled on the specified management interface address: HTTP, HTTPS, Telnet, Secure Shell (SSH), and/or ping.
Permitted IPs	Enter the list of IP addresses from which firewall management is allowed. When using this option for Panorama, you will need to make sure that each managed device has it's IP address added, otherwise it will not be able to connect send logs to Panorama or receive configuration updates.

Table 1. Management Settings (Continued)

ltem	Description
Logging and Reporting Settings	
Log Storage	Specify the percentage of space allocated to each log type on the hard disk.
	When you change a percent value, the associated disk allocation changes automatically. If the total of all the values exceeds 100%, a message appears on the page in red, and an error message is presented when you attempt to save the settings. If this occurs, readjust the percentages so the total is within the 100% limit.
	Click OK to save settings and Restore Defaults to restore all of the default settings.
	Note: When a log reaches its maximum size, it starts to be overwritten beginning with the oldest entries. If you resize an existing log to be smaller than its current size, the firewall starts immediately to cut down the log when you commit the changes, with the oldest logs removed first.
Max Rows in User Activity Report	Enter the maximum number of rows that is supported for the detailed user activity reports (1-1048576, default 65535).
Max Rows in CSV Export	Enter the maximum number of rows that will appear in the CSV reports generated from the Export to CSV icon in the traffic logs view (range 1-1048576, default 65535).
Number of Versions for Config Audit	Enter the number of configuration audit versions to save before discarding the oldest ones (default 100).
Number of Versions for Config Backups	(Panorama only) Enter the number of configuration backups to save before discarding the oldest ones (default 100).
Average Browse Time (sec)	Configure this variable to adjust how browse time is calculated in the "User Activity Report".
	The calculation will ignore sites categorized as web advertisements and content delivery networks. The browse time calculation is based on container pages logged in the URL filtering logs. Container pages are used as the basis for this calculation because many sites load content from external sites that should not be considered. For more information on the container page, refer to "Container Pages" on page 45.
	The average browse time setting is the average time that the admin thinks it should take a user to browse a web page. Any request made after the average browse time has elapsed will be considered a new browsing activity. The calculation will ignore any new web pages that are loaded between the time of the first request (start time) and the average browse time. This behavior was designed to exclude any external sites that are loaded within the web page of interest.
	Example: If the average browse time setting is 2 minutes and a user opens a web page and views that page for 5 minutes, the browse time for that page will still be 2 minutes. This is done because there is no way to determine how long a user views a given page.
	(Range 0-300 seconds, default 60 seconds)

Table 1. Management Settings (Continued)

ltem	Description
Page Load Threshold (sec)	Configure this variable to adjust how browse time is calculated in the "User Activity Report".
	This option allows you to adjust the assumed time it takes for page elements to load on the page. Any request that occurs between the first page load and the page load threshold is assumed to be elements of the page. Any requests that occur outside of the page load threshold is assumed to be the user clicking a link within the page.
Con III of a second la	(Kange 0-60 seconds, default 20 seconds)
Send Hostname In Syslog	When this option is set, syslog messages will contain the hostname of the firewall device in their header.
Stop Traffic when LogDb full	Select the check box if you want traffic through the firewall to stop when the log database is full (default off).
Enable Log on High DP Load	Select this check box if you would like a system log entry generated if the device is under severe load (default off).
Buffered log forwarding from device (option in the Panorama > Set up tab)	Allows the firewall to buffer log entries on the device's hard disk (local storage) when it loses connectivity to Panorama. When the connection to Panorama is restored, the log entries are forwarded to Panorama; the disk space available for buffering depends on the log storage quota for the platform and the volume of logs that are pending roll over. If the available space is consumed, the oldest entries are deleted to allow logging of new events. Enabled by default.
Get Only New Logs on Convert to Primary (option in the Panorama > Set up tab)	This option is only applicable when Panorama writes logs to a Network File Share (NFS). With NFS logging, only the <i>primary</i> Panorama is mounted to the NFS. Therefore, the devices send logs to the <i>active primary</i> Panorama only.
	This option allows an administrator to configure the managed devices to only send newly generated logs to Panorama when an HA failover occurs and the secondary Panorama resumes logging to the NFS (after it is promoted as primary).
	This behavior is typically enabled to prevent the devices from sending a large volume of buffered logs when connectivity to Panorama is restored after a significant period of time.
Only Active Primary Logs to Local Disk (in Panorama > Set up tab only)	Allows you to configure only the active primary Panorama to save logs to the local disk.
	This option is valid for a Panorama virtual machine with a virtual disk and to the M-100 appliance in Panorama mode.
Minimum Password Complexity	
Enabled	Enable minimum password requirements for local accounts. With this feature, you can ensure that local administrator accounts on the firewall will adhere to a defined set of password requirements.
	You can also create a password profile with a subset of these options that will override these settings and can be applied to specific accounts. For more information, refer to "Defining Password Profiles" on page 59 and refer to "Username and Password Requirements" on page 57 for information on valid characters that can be used for accounts.

Table 1. Management Settings (Continued)

ltem	Description
	Note: The maximum password length that can be entered is 31 characters. When setting requirements, make sure you do not create a combination that will not be accepted. Example, you would not be able to set a requirement of 10 uppercase, 10 lower case, 10 numbers, and 10 special characters since that would exceed the maximum length of 31.
	Note: If you have High Availability (HA) configured, always use the primary device when configuring password complexity options and commit soon after making changes.
Minimum Length	Require minimum length from 1-15 characters.
Minimum Uppercase Letters	Require a minimum number of uppercase letters from 0-15 characters.
Minimum Lowercase Letters	Require a minimum number of lowercase letters from 0-15 characters.
Minimum Numeric Letters	Require a minimum number of numeric letters from 0-15 numbers.
Minimum Special Characters	Require a minimum number of special characters (non-alphanumeric) from 0-15 characters.
Block Repeated Characters	Do no allow repeated characters based on the specified value. Example, if the value is set to 4, the password test2222 would not be accepted, but test2222 would be accepted (range 2-15).
Block Username Inclusion (including reversed)	Select this check box to prevent the account username (or reversed version of the name) from being used in the password.
New Password Differs By Characters	When administrators change their passwords, the characters must differ by the specified value.
Require Password Change on First Login	Select this check box to prompt the administrators to change their passwords the first time they log in to the device.
Prevent Password Reuse Limit	Require that a previous password is not reused based on the specified count. Example, if the value is set to 4, you could not reuse the any of your last 4 passwords (range 0-50).
Block Password Change Period (days)	User cannot change their passwords until the specified number of days has been reached (range 0-365 days).
Required Password Change Period (days)	Require that administrators change their password on a regular basis specified a by the number of days set, ranging from 0-365 days. Example, if the value is set to 90, administrators will be prompted to change their password every 90 days.
	You can also set an expiration warning from 0-30 days and specify a grace period.
Expiration Warning Period (days)	If a required password change period is set, this setting can be used to prompt the user to change their password at each log in as the forced password change date approaches (range 0-30 days).

Table 1. Management Settings (Continued)
ltem	Description
Allowed expired admin login (count)	Allow the administrator to log in the specified number of times after the account has expired. Example, if the value is set to 3 and their account has expired, they can log in 3 more times before their account is locked out (range 0-3 logins).
Post Expiration Grace Period (days)	Allow the administrator to log in the specified number of days after the account has expired (range 0-30 days).

Table 1. Management Settings (Continued)

Defining Operations Settings

Device > Setup > Operations

When you change a configuration setting and click **OK**, the current "candidate" configuration is updated, not the active configuration. Clicking **Commit** at the top of the page applies the candidate configuration to the active configuration, which activates all configuration changes since the last commit.

This method allows you to review the configuration before activating it. Activating multiple changes simultaneously helps avoid invalid configuration states that can occur when changes are applied in real-time.

You can save and roll back (restore) the candidate configuration as often as needed and also load, validate, import, and export configurations. Pressing **Save** creates a copy of the current candidate configuration, whereas choosing **Commit** updates the active configuration with the contents of the candidate configuration.



Note: It is a good idea to periodically save the configuration settings you have entered by clicking the *Save* link in the upper-right corner of the screen.

To manage configurations, select the appropriate configuration management functions, as described in the following table.

Function	Description
Configuration Management	
Validate candidate config	Checks the candidate configuration for errors.
Revert to last saved config	Restores the last saved candidate configuration from the local drive. The current candidate configuration is overwritten. An error occurs if the candidate configuration has not been saved.
Revert to running config	Restores the last running configuration. The current running configuration is overridden.
Save named configuration snapshot	Saves the candidate configuration to a file. Enter a file name or select an existing file to be overwritten. Note that the current active configuration file (<i>running-config.xml</i>) cannot be overwritten.

Table 2. Configuration Management Functions

Function	Description
Save candidate config	Saves the candidate configuration in flash memory (same as clicking Save at the top of the page).
Load named configuration snapshot	Loads a candidate configuration from the active configuration (<i>running-config.xml</i>) or from a previously imported or saved configuration. Select the configuration file to be loaded. The current candidate configuration is overwritten.
Load configuration version	Loads a specified version of the configuration.
Export named configuration snapshot	Exports the active configuration (<i>running-config.xml</i>) or a previously saved or imported configuration. Select the configuration file to be exported. You can open the file and/or save it in any network location.
Export configuration version	Exports a specified version of the configuration.
Export device state	This feature is used to export the configuration and dynamic information from a firewall that is configured as a GlobalProtect Portal with the large scale VPN feature enabled. If the Portal experiences a failure, the export file can be imported to restore the Portal's configuration and dynamic information.
	The export contains a list of all satellite devices managed by the Portal, the running configuration at the time of the export, and all certificate information (Root CA, Server, and Satellite certificates).
	Important: You must manually run the device state export or create a scheduled XML API script to export the file to a remote server. This should be done on a regular basis since satellite certificates may change often.
	To create the device state file from the CLI, from configuration mode run save device state. The file will be named device_state_cfg.tgz and is stored in /opt/pancfg/ mgmt/device-state. The operational command to export the device state file is scp export device-state (you can also use tftp export device-state).
	For information on using the XML API, refer to the document "PAN-OS XML-Based Rest API Usage Guide" at <i>https://live.paloaltonetworks.com/community/documentation</i> .
	Refer to "GlobalProtect Large Scale VPN Deployment" on page 323.
Import named config snapshot	Imports a configuration file from any network location. Click Browse and select the configuration file to be imported.
Import device state	Import the device state information that was exported using the Export device state option. This includes the current running config, Panorama templates, and shared policies. If the device is a Global Protect Portal, the export includes the Certificate Authority (CA) information and the list of satellite devices and their authentication information.
Device Operations	

 Table 2. Configuration Management Functions (Continued)

Function	Description
Reboot Device	To restart the firewall, click Reboot Device . You are logged out and the PAN-OS software and active configuration are reloaded. Existing sessions will also be closed and logged and a system log entry will be created that will show the administrator name that initiated the shutdown. Any configuration changes that have not been saved or committed are lost (refer to "Defining Operations Settings" on page 37).
	Note: If the web interface is not available, use the CLI command request restart system. Refer to the PAN-OS Command Line Interface Reference Guide for details.
Shutdown Device	To perform a graceful shutdown of the firewall, click Shutdown Device and then click Yes on the confirmation prompt. Any configuration changes that have not been saved or committed are lost. All administrators will be logged off and the following processes will occur:
	 All login sessions will be logged off.
	• Interfaces will be disabled.
	• All system processes will be stopped.
	• Existing sessions will be closed and logged.
	• System Logs will be created that will show the administrator name who initiated the shutdown. If this log entry cannot be written, a warning will appear and the system will not shutdown.
	• Disk drives will be cleanly unmounted and the device will powered off.
	You need to unplug the power source and plug it back in before you can power on the device.
	Note: If the web interface is not available, use the CLI command request shutdown system. Refer to the PAN-OS Command Line Interface Reference Guide for details.
Restart Data Plane	To restart the data functions of the firewall without rebooting, click Restart Dataplane . <i>This option is not available on the PA-200</i> .
	Note: If the web interface is not available, use the CLI command request restart dataplane. Refer to the PAN-OS Command Line Interface Reference Guide for details.

Table 2. Configuration Management Functions (Continued)

Function	Description
Miscellaneous	
Custom Logos	Use this option to customize any of the following:
	 Login screen background image
	• Main UI (User Interface) header image
	 PDF report title page image. Refer to "Managing PDF Summary Reports" on page 270.
	• PDF report footer image
	Click 🍐 to upload an image file, 🔎 to preview, or 😑 to remove a previously-uploaded image.
	Note the following:
	 Supported file types are png, gif, and jpg.
	Note: Image files that contain an alpha channel are not supported and when used in PDF reports, the reports will not be generated properly. You may need to contact the illustrator who created the image to remove alpha channels in the image or make sure the graphics software you are using does not save files with the alpha channel feature.
	• To return to the default logo, remove your entry and commit.
	• The maximum image size for any logo image is 128 KB.
	• For the login screen and main user interface options, when you click , the image is shown as it will be displayed. If necessary, the image is cropped to fit. For the PDF reports, the images are auto-resized to fit without cropping. In all cases, the preview shows the recommended image dimensions.
	For information on generating PDF reports, refer to "Managing PDF Summary Reports" on page 270.
SNMP Setup	Specify SNMP parameters. Refer to "SNMP" on page 47.
Statistics Service Setup	Specify settings for the statistics service. Refer to "Statistics Service" on page 48.

Table 2. Configuration Management Functions (Continued)



Note: When you click *Commit* or enter a *commit* CLI command, all changes made through the web interface and the CLI since the last commit are activated. To avoid possible conflicts, use the transaction locking functions as described in "Locking Transactions" on page 27.

Defining Services Settings

Device > Setup > Services

Use the **Services** tab to define settings for Domain Name System (DNS), Network Time Protocol (NTP), update servers, proxy servers, and service route configuration.

Table 3. Services Settings

Function	Description
DNS	Select the type of DNS service. This setting is used for all DNS queries initiated by the firewall in support of FQDN address objects, logging, and device management. Options include:
	 Primary and secondary DNS servers for domain name resolution
	• DNS proxy that has been configured on the firewall
Primary DNS Server	Enter the IP address of the primary DNS server. The server is used for DNS queries from the firewall, for example, to find the update server, to resolve DNS entries in logs, or for FDQN-based address objects.
Secondary DNS Server	Enter the IP address of a secondary DNS server to use if the primary server is unavailable (optional).
Primary NTP Server	Enter the IP address or host name of the primary NTP server, if any. If you do not use NTP servers, you can set the device time manually.
Secondary NTP Server	Enter the IP address or host name of secondary NTP servers to use if the primary server is unavailable (optional).
Update Server	This setting represents the IP address or host name of the server used to download updates from Palo Alto Networks. The current value is updates.paloaltonetworks.com . Do not change the server name unless instructed by technical support.
Proxy Server	
Server	If the device needs to use a proxy server to reach Palo Alto Networks update services, enter the IP address or host name of the server.
Port	Enter the port for the proxy server.
User	Enter the user name to access the server.
Password/Confirm Password	Enter and confirm the password for the user to access the proxy server.

Function	Description
Service Route Configuration	Specify how the firewall will communicate with other servers/devices for services communication, such as DNS, Email, Palo Alto Updates, and NTP. You can select "Use Management Interface for all" to use the built-in management port (MGT) for all communications, or you can define a specific source IP address to use different interfaces for granular control of each service. For example, you could configure a specific source IP of an interface for all email communication between the firewall and an email server and use a different source IP/interface for Palo Alto Updates
	Click Service Route Configuration and configure the following:
	• Use Management Interface for all—This option will force all firewall service communications with external servers through the management interface (MGT). If this option is selected, you will need to configure the MGT interface to allow communications between the firewall and the servers/devices that provide services. To configure the MGT interface, navigate to Device > Setup > Management tab and edit "Management Interface Settings".
	• Select—Choose this option to configure granular control for service communication. A table will appear that shows a list of available services and a drop-down to select a source address. Select the desired service and then select the source from the Source Address drop-down list. The "Source Address" is the IP address assigned to an interface that will be the source for the service traffic. You do not have to define the destination address since the destination is configured when configuring the given service. For example, when you define your DNS servers from the Device > Setup > Services tab and Services, that will set the destination for DNS queries.

 Table 3.
 Services Settings (Continued)

Function	Description
Service Route Configuration (Continued)	• Destination and Source Address fields—If a service that you want to route is not listed in the Service column, you can use the Destination and Source Address fields to define routes that will be used by other services. Services not listed include items such as Kerberos, LDAP, and Panorama log collector communications. You do not need to enter the subnet for the destination address.
	In multi-tenant environments, destination IP-based service routes will be required where common services require different source address. For example, if two tenants need to use RADIUS.
	It is important that routing and policies are setup properly for the interface that will be used to route the service. For example, if you want to route Kerberos authentication requests on an interface other than the MGT port, you need to configure the Destination and Source Address in the right section of the Service Route Configuration window since Kerberos is not listed in the default Service column. As an example, you could have a source IP address 192.168.2.1 on Ethernet1/3 and then a destination for a Kerberos server of 10.0.0.240. You will need to add Ethernet1/3 to an existing virtual router with a default route, or you can create a new virtual router from Network > Virtual Routers and add static routes as needed. This will ensure that all traffic on the interface will be routed through the virtual router to reach the appropriate destinations. In this case, the destination address is 10.0.240.
	The CLI output for the Destination and Source Address would look like the following:
	PA-200-Test# show route
	destination {
	10.0.240 {
	source address 192.168.2.1/24
	} With this configuration, all traffic on interface Ethernet1/3 will use the default route defined in the virtual router and will be sent to 10.0.0.240.

Defining Content ID Settings

► Device > Setup > Content-ID

Use the **Content-ID** tab to define settings for URL filtering, data protection, and container pages.

Table 4. Content ID Settings

Function	Description
URL Filtering	

Function	Description
Dynamic URL Cache Timeout	Click Edit and enter the timeout (in hours). This value is used in dynamic URL filtering to determine the length of time an entry remains in the cache after it is returned from the URL filtering service. This option is applicable to URL filtering using the BrightCloud database only. For information on URL filtering, refer to "URL Filtering Profiles" on page 217.
URL Continue Timeout	Specify the interval following a user's "continue" action before the user must press continue again for URLs in the same category (range 1 - 86400 minutes, default 15 minutes).
URL Admin Override Timeout	Specify the interval after the user enters the admin override password before the user must re-enter the admin override password for URLs in the same category (range 1 - 86400 minutes, default 900 minutes).
URL Admin Lockout Timeout	Specify the period of time that a user is locked out from attempting to use the URL Admin Override password following three unsuccessful attempts (1 - 86400 minutes, default 1800 minutes).
x-forwarded-for	Include the X-Forwarded-For header that includes the source IP address. When this option is selected, the firewall examines the HTTP headers for the X-Forwarded-For header, which a proxy can use to store the original user's source IP address.
	The system takes the value and places Src: x.x.x. into the Source User field of the URL logs (where <i>x.x.x.x</i> is the IP address that is read from the header).
Strip-x-forwarded-for	Remove the X-Forwarded-For header that includes the source IP address. When this option is selected, the firewall zeros out the header value before forwarding the request, and the forwarded packets do not contain internal source IP information.
Allow Forwarding of Decrypted Content	Select the check box to allows the firewall to forward decrypted content to an outside service. For example, when this option is set the firewall can send decrypted content to WildFire for analysis. For multi-VSYS configurations, this option is per VSYS.
URL Admin Override	
Settings for URL Admin Override	Specify the settings that are used when a page is blocked by the URL filtering profile and the Override action is specified. Refer to "URL Filtering Profiles" on page 217.
	Click Add and configure the following settings for each virtual system that you want to configure for URL admin override.
	• Location—Select the virtual system from the drop-down list (multi- VSYS devices only).
	• Password/Confirm Password —Enter the password that the user must enter to override the block page.
	• Server Certificate—Select the server certificate to be used with SSL communications when redirecting through the specified server.
	• Mode —Determines whether the block page is delivered transparently (it appears to originate at the blocked website) or redirected to the user to the specified server. If you choose Redirect , enter the IP address for redirection.
	Click 🔣 to delete an entry.

Table 4. Content ID Settings (Continued)

Function	Description
Manage Data Protection	Add additional protection for access to logs that may contain sensitive information, such as credit card numbers or social security numbers.
	Click Manage Data Protection and configure the following:
	• To set a new password if one has not already been set, click Set Pass-word . Enter and confirm the password.
	• To change the password, click Change Password . Enter the old password, and enter and confirm the new password.
	• To delete the password and the data that has been protected, click Delete Password .
Container Pages	Use these settings to specify the types of URLs that the firewall will track or log based on content type, such as application/pdf, application/ soap+xml, application/xhtml+, text/html, text/plain, and text/xml. Container pages are set per virtual system, which you select from the Location drop-down list. If a virtual system does not have an explicit container page defined, the default content types are used.
	Click Add and enter or select a content type.
	Adding new content types for a virtual system overrides the default list of content types. If there are no content types associated with a virtual system, the default list of content types is used.

Table 4. Content ID Settings (Continued)

Defining Session Settings

► Device > Setup > Session

The **Session** tab allows you to configure session age-out times and global session-related settings such as firewalling IPv6 traffic and rematching security policy to existing sessions when the policy changes.

Field	Description		
Session Settings			
Rematch Sessions	Click Edit and select the check box Rematch all sessions on config policy change .		
	For example, assume that Telnet was previously allowed and then changed to Deny in the last commit. The default behavior is for any Telnet sessions that were started before the commit to be rematched and blocked.		
ICMPv6 Token Bucket Size	Enter the bucket size for rate limiting of ICMPv6 error messages. The token bucket size is a parameter of the token bucket algorithm that controls how bursty the ICMPv6 error packets can be (range 10-65535 packets, default 100).		
ICMPv6 Error Packet Rate	Enter the average number of ICMPv6 error packets per second allowed globally (range 10-65535 packets/sec, default 100). This value applies to all interfaces.		
Jumbo Frame Jumbo Frame MTU	Select to enable jumbo frame support. The Default value for Jumbo frames is 9192. (range 512 - 9216 bytes). Refer to the spec sheet for your firewall model, available at <i>http://www.paloaltonetworks.com</i> .		

Table 5.Session Settings

Field	Description		
Enable IPv6 Firewalling	To enable firewall capabilities for IPv6, click Edit and select the IPv6 Firewalling check box.		
	All IPv6-based configurations are ignored if IPv6 is not enabled.		
NAT64 IPv6 Minimum Network MTU	Allows you to change the global MTU setting for IPv6 translated traffic (default 1280). The default is based on the standard minimum MTU for IPv6 traffic.		
Accelerated Aging	Allows for the accelerated aging-out of idle sessions.		
	Select the check box to enable accelerated aging and specify the threshold (%) and scaling factor.		
	When the session table reaches the Accelerated Aging Threshold (% full), the Accelerated Aging Scaling Factor is applied to the aging calculations for all sessions. The session's idle time is calculated as the actual idle time multiplied by the scaling factor. For example, if a scaling factor of 10 is used, a session that would normally time out after 3600 seconds would instead time out 10 times faster, so it would time out in 360 seconds.		
Session Timeouts			
Timeouts	Specify timeouts in seconds for each of the categories. Ranges and defaults are listed.		
Session Features			
Decryption Certificate Revocation Settings	Select to configure the certificate management options for the firewall.		
Enable	Select the check box to use Certificate Revocation Lists (CRL) to check the validity of SSL certificates.		
	Each trusted certificate authority (CA) maintains CRL to determine if an SSL certificate is valid (not revoked) for SSL decryption. The Online Certificate Status Protocol (OCSP) can also be used to dynamically check the revocation status of a certificate. For more information on SSL decryption, refer to "Decryption Policies" on page 202.		
Receive Timeout	Specify the interval after which the CRL request times out and the status is determined to be unknown (1-60 seconds).		
Enable OCSP	Select the check box to use OCSP to check the validity of SSL certificates.		
Receive Timeout	Specify the interval after which the OCSP requests times out and the status is determined to be unknown (1-60 seconds).		
Block Session With Unknown Certificate Status	Select the check box if you want to block certificates that cannot be validated.		
Block Session On Certificate Status Check Timeout	Select the check box if you want to block certificates when the request for certificate information times out.		
Certificate Status Timeout	Specify the interval after which certificate status requests time out (1-60 seconds).		

Table 5.	Session	Settings	(Continued)
----------	---------	----------	-------------

SNMP

Device > Setup > Operations

Use this page to define access to SNMP Management Information Bases (MIBs) for SNMPv2c and SNMPv3. Click **SNMP Setup** on the **Setup** page, and specify the following settings.

A MIB module defines all SNMP traps generated by the system. Each event log in the system is defined as an independent SNMP trap with an Object ID (OID) of its own, and individual fields in an event log are defined as a variable binding (varbind) list.

Field	Description			
Physical Location	Specify the physical location of the firewall.			
Contact	Enter the name or email address of the person responsible for maintaining the firewall. This setting is reported in the standard system information MIB.			
Use Specific Trap Definitions	Select the check box to use a unique OID for each SNMP trap based on the event type (default is selected).			
Version	Select the SNMP version (V2c or V3). This setting controls access to the MIB information. By default, V2c is selected with the "public" community string.			
	For V2c, configure the following setting:			
	• SNMP Community String —Enter the SNMP community string for firewall access (default public).			
	For V3, configure the following settings:			
	• Views—Click Add and configure the following settings:			
	 Name—Specify a name for a group of views. 			
	 View—Specify a name for a view. 			
	- OID —Specify the object identifier (OID) (for example, 1.2.3.4).			
	 Option—Choose whether the OID is to be included or excluded from the view. 			
	 Mask—Specify a mask value for a filter on the OID in hexadecimal format (for example, 0xf0). 			
	• Users—Click Add and configure the following settings:			
	– Users —Specify a user name.			
	 View—Specify the group of views for the user. 			
	 Auth Password—Specify the user's authentication password (minimum 8 characters, maximum of 256 characters, and no character restrictions). All characters allowed). Only Secure Hash Algorithm (SHA) is supported. 			
	 Priv Password—Specify the user's encryption password (minimum 8 characters, maximum of 256 characters, and no character restrictions). Only Advanced Encryption Standard (AES) is supported. 			

Table 6. SNMP Setup

Statistics Service

Device > Setup > Operations

The **Statistics Service** feature allows the firewall to send anonymous application, threat, and crash information to the Palo Alto Networks research team. The information collected enables the research team to continually improve the effectiveness of Palo Alto Networks products based on real-world information. This service is disabled by default and once enabled, information will be uploaded every 4 hours.

You can allow the firewall to send any of the following types of information:

- Application and Threat Reports
- Unknown Application Reports
- URL Reports
- Device traces for crashes

To view a sample of the content for a statistical report to be sent, click the report icon is . The **Report Sample** tab opens to display the report code. To view a report, click the check box next to the desired report, then click the **Report Sample** tab.

Comparing Configuration Files

Device > Config Audit

You can view and compare configuration files by using the **Config Audit** page. From the drop-down lists, select the configurations to compare. Select the number of lines that you want to include for context, and click **Go**.

The system presents the configurations and highlights the differences, as in the following figure.

The page also includes *«* and *»* buttons adjacent to the drop-down lists, which are enabled when comparing two consecutive configuration versions. Click *«* to change the configurations being compared to the previous set of stored configurations, and click to *»* to change the configurations being compared to the next set of stored configurations.



Figure 1. Configuration Comparison

Panorama automatically saves all of the configuration files that are committed on each managed firewall, whether the changes are made through the Panorama interface or locally on the firewall.

Installing a License

Device > Licenses

When you purchase a subscription from Palo Alto Networks, you receive an authorization code to activate one or more license keys.

To activate a URL filtering license, you must install the license, download the database, and click **Activate**.

The following functions are available on the Licenses page:

- To enable licenses for URL filtering, click **Activate**. If you are using PAN-DB for URL Filtering, you will need to click **Download** to retrieve the initial seed database first and then click Activate. You can also run the CLI request url-filtering download paloaltonetworks region <region name>.
- To enable purchased subscriptions that require an authorization code and have been activated on the support portal, click **Retrieve license keys from license server**.
- To enable purchased subscriptions that require an authorization code and have not been previously activated on the support portal, click **Activate feature using authorization code**. Enter your authorization code, and click **OK**.
- If the firewall does not have connectivity to the license server and you want to upload license keys manually, follow these steps:
 - a. Obtain a file of license keys from *http://support.paloaltonetworks.com*.
 - b. Save the license key file locally.
 - c. Click Manually upload license key, click Browse and select the file, and click OK.

Important items to consider when installing a license

If you are unable to activate the URL filter using the web interface, CLI commands are available. Refer to the *PAN-OS Command Line Interface Reference Guide* for more information.

Upgrading/Downgrading the PAN-OS Software

► Device > Software

To upgrade to a new release of the PAN-OS software, you can view the latest versions of the PAN-OS software available from Palo Alto Networks, read the release notes for each version, and then select the release you want to download and install (a support license is required).

Perform any of the following functions on the **Software** page:

- Click Check Now to retrieve the latest software releases available from Palo Alto Networks.
- Click Release Notes to view a description of the changes in a release and to view the migration path to install the software.
 You cannot skip a feature release and must have a base image downloaded before upgrading from one feature release to a maintenance release in a higher feature release. This is due to the fact that the maintenance release does not contain the full software

build, only changes made since the base image release, so the base image is needed for the upgrade. For example, if you upgrade from 4.0.12 to 4.1.7, you need to download the 4.1.0 base image (not install), so the 4.1.7 maintenance release upgrade can access the 4.1.0 base image files. If you want to upgrade to from one release to two higher feature releases, you need to upgrade to each feature release. For example, if you want to upgrade from 4.0 to 5.0, you have to upgrade from 4.0 to 4.1, then 4.1 to 5.0.

You can delete the base image files when you complete the upgrade, but that is not recommended since the base image would be needed when you upgrade to the next maintenance release. You would only want to delete the base image for older releases that will not need upgrading. For example, if you are running 4.1, you probably do not need the base images for 3.1 and 4.0 unless you plan on downgrading to those versions.

• Click **Download** to install a new release from the download site. When the download is complete, a checkmark is displayed in the **Downloaded** column. To install a downloaded release, click **Install** next to the release.

During installation, you are asked whether to reboot when installation is complete. When the installation is complete, you will be logged out while the firewall is restarted. The firewall will be rebooted, if that option was selected.

When upgrading to a feature release (where the first or second digit in the PAN-OS version changes, e.g. 4.0 to 4.1. or 4.1 to 5.0), you will see a message when you click install that indicates that you are about to perform a feature release upgrade. You should make sure you back up your current configuration since a feature release may migrate certain configurations to accommodate new features. Refer to "Downgrading PAN-OS Software" on page 53.

- Click **Upload** to install a release that you previously stored on your PC. Browse to select the software package, and click **Install from File**. Choose the file that you just selected from the drop-down list, and click **OK** to install the image.
- Click the Delete icon 🕅 to delete an outdated release.

Items to note when upgrading the PAN-OS software

- When upgrading from an earlier PAN-OS version, follow the recommended path to reach the latest release, as described in the release notes.
- When upgrading a High Availability (HA) pair to a new feature release (where the first or second digit in the PAN-OS version changes, e.g. 4.0 to 4.1. or 4.1 to 5.0), the configuration may be migrated to accommodate new features. If session synchronization is enabled, sessions will not be synchronized if one device in the cluster is at a different PAN-OS feature release.
- The date and time settings on the firewall must be current. PAN-OS software is digitally signed and the signature is checked by the device prior to installing a new version. If the date setting on the firewall is not current, the device may perceive the software signature to be erroneously in the future and will display the message Decrypt failed: GnuPG edit non-zero, with code 171072 Failed to load into PAN software manager.

Upgrading PAN-OS in a High Availability Configuration

This section provides information on upgrading the PAN-OS software on a high availability (HA) configuration in active/passive and active/active modes.

If session synchronization is enabled, sessions will not be synchronized if one device in the cluster is at a different PAN-OS feature release. If session continuity is required, you must temporarily permit non-syn-tcp while the session table is rebuilt.

For information on High Availability (HA) configurations, refer to "High Availability" on page 93. To view the general status of your HA configuration, go to the **Dashboard** tab and view the **High Availability** widget. If the widget is not displayed, click the Widgets link drop-down and select **System > High Availability**.

The following steps assume that you are performing a new feature release upgrade. The same procedures can be used for a maintenance release, but in that case if session synchronization is enabled, there is no interruption to that process while both devices are functional. Also, the higher revision maintenance release in the pair will not go non-functional as it does with a feature release.

To upgrade an HA pair to a new feature release:



Note: In the following steps, active/passive and active/active applies, the terminology is just different. In active/passive you have an active and a passive device. In active/active you have an active-primary and active-secondary and both devices are passing traffic when the cluster is functional.

- 1. Read "Upgrading/Downgrading the PAN-OS Software" on page 50 for general information and important notes on the upgrade process.
- Save a backup of the current configuration file on both devices by navigating to Device > Setup > Operations tab and select Export named configuration snapshot, select running-config.xml and then click OK to save the configuration file to your computer.
- 3. Upgrade the passive/active-secondary firewall (device B) to the new PAN-OS feature release and then reboot to complete the installation. After the reboot, the device will not be functional until device A (active/active-primary) is suspended. If session synchronization is enabled, you can run the operational command set session tcp-reject-non-syn no. This will rebuild the session table and sessions prior to the upgrade will continue. Once both devices are at the same feature release, run set session tcp-reject-non-syn yes to enable this option. In active/passive mode, if the preemptive option is configured, the current passive device will revert to active when state synchronization is complete.
- 4. Suspend the active/active-primary firewall (device A), which will force the passive/ active-secondary (device B) firewall to become functional. At this point, session synchronization will stop due to the fact that both devices are not on the same PAN-OS feature release. Make sure device B is functional and passing traffic before continuing the upgrade on device A.
- 5. Upgrade "device A" to the new PAN-OS release and then reboot the device to complete the installation. When device A comes back up, all HA functions should resume, including session synchronization. If you are allowing non-syn-tcp as indicated in Step 3, you can revert back by running set session tcp-reject-non-syn yes.
- 6. Verify that the active device is passing traffic by viewing the Monitor > Session Browser, or by running show session all from the CLI. You can also check the state of HA on the device by running show high-availability all | match reason. If this is an active/active configuration, check that both devices are passing traffic. To check session synchronization run show high-availability interface ha2. In the

52 • Device Management

"Hardware Interface counters read from CPU" table check that counters are increasing. In an active/passive configuration, only the active device will show packets transmitted and the passive device will only show packets received. In active/active, you will see packets received and packets transmitted on both devices.

Downgrading PAN-OS Software

If you need to downgrade the PAN-OS software, you will need to follow different steps based on the PAN-OS release type (feature or maintenance). In a feature release (where the first or second digit in the PAN-OS version changes, e.g. 4.0 to 4.1. or 4.1 to 5.0), the configuration may be migrated to accommodate new features, so you should not downgrade unless you also restore the configuration for the previous release. Maintenance releases can be downgraded without having to worry about restoring the configuration.

- "Maintenance Release Downgrade" in the next section
- "Feature release Downgrade" on page 54



CAUTION: It is recommended that you downgrade into a configuration that matches the software version. Unmatched software and configurations can result in failed downgrades or even force the system into maintenance mode. This only applies to a downgrade from one feature release to another, not maintenance releases.

If you have a problem with a downgrade, you may need to enter maintenance mode and reset the device to factory default and then restore the configuration from the original config file that was exported prior to the upgrade.

Maintenance Release Downgrade

A maintenance release is where the third digit in the release numbers changes (4.1.4 to 4.1.5 or 5.0.0 to 5.0.1). When you upgrade from one maintenance release to another, there are no feature changes, so there is no need to restore the configuration during the downgrade.

To downgrade to an earlier maintenance release:

- Save a backup of the current configuration file by navigating to the Device > Setup >
 Operations tab and select Export named configuration snapshot, select running config.xml and then click OK to save the configuration file. This backup can be used to
 restore the configuration if you have problems with the downgrade and you need to do a
 factory reset.
- 2. Navigate to **Device > Software** and you will see the software page that lists all PAN-OS versions that can be downloaded, or that have already been downloaded.
- 3. To downgrade to an older maintenance release, click **Install** in the **Action** column for the desired release. If the version you want to use shows **Download**, click the **Download** link to retrieve the software package and then click **Install**.
- 4. After PAN-OS has been downgraded, click **OK** to reboot the device to activate the new version.

Feature release Downgrade

A feature release is where the first or second digit in the release numbers changes (4.0 to 4.1, or 4.1 to 5.0). When you upgrade from one feature release to another, the configuration may be changed to accommodate the new features. Due to this fact, you will need to restore your configuration backup that was created before the device was upgraded to a feature release. As of PAN-OS 4.1, this configuration backup is automatically created when upgrading to a feature release.

To downgrade to an earlier feature release:



Note: It is important to note that this procedure will restore your device to the configuration that was in place before the upgrade to a feature release. Any changes made since that time will be lost, so it is important to back up your current configuration in case you want to restore those changes when you return to the newer release.

- Save a backup of the current configuration file by navigating to the Device > Setup >
 Operations tab and select Export named configuration snapshot, select running config.xml and then click OK to save the configuration file. This backup can be used to
 restore the configuration if you have problems with the downgrade and you need to do a
 factory reset.
- 2. To downgrade to an older feature release, you must locate the previous feature release by navigating to **Device > Software** and browsing to the page that contains the release. See Figure 2.
- 3. Click **Install** in the **Action** column for the feature release. If the version you want to use shows **Download**, click the **Download** link to retrieve the software package and then click **Install**.
- 4. You will see a prompt that allows you to select a configuration that will be used after the device reboots. In most cases, since this is a feature release downgrade, you want to select the auto saved configuration that was created when the device was upgraded to the next feature release. For example, if you are running PAN-OS 5.0 and want to downgrade to

PAN-OS 4.1, click **Install** in the **Action** column, then in the **Select a Config File for Downgrading** drop-down, select *autosave-4.1.0* as shown in Figure 2.

Version	Size	Release Date	Downlo	Currently Installed	Action	
4.1.1-h2	145 MB	2011/12/27 13:52:54			Download	Release Notes
4.1.1-h1	145 MB	2011/12/12 16:45:19			Download	Release Notes
(4.1.0)	248 MB	2011/10/30 20:00:23	~		Install	Release Notes
4.1.0-c1		2011/10/28 18:05:57			Download	Release Notes
4.1.0-c1		2011/10/27 19:23:04			Download	Release Notes
4.1.0-c1		2011/10/24 19:48:34			Download	Release Notes
4.1.0-c1		² Select A Config File fo	r Downgra	ding		0
4.1.0-c1		2				
4.1.0-c1		2 Name (autosave-4.1.0)				
4.1.0-c1		2 If the config version is incompatible with the SW version,				
4.1.0-c1		2 the sys	stem downgra	de may fail.		
∢ 4.1.0-c1		2				
4.1.0-c1		2			ОК	Close
4.1.0-c1		2011/10/12 08:00:25			Download	Release Notes

Figure 2. Downgrade to an older feature release

5. After PAN-OS has been installed, click **OK** to reboot the device to activate the new version. The selected feature release of PAN-OS and the configuration will then be activated.



Note: For releases prior to 4.1, you may need to do a factory reset and restore the device. For example, downgrading from 4.0 to 3.1.

Updating Threat and Application Definitions

Device > Dynamic Updates

Palo Alto Networks periodically posts updates with new or revised application definitions, information on new security threats, such as antivirus signatures (threat prevention license required), URL filtering criteria, updates to GlobalProtect data, and WildFire signatures (WildFire subscription required). You can view the latest updates, read the release notes for each update, and then select the update you want to download and install. You can also revert to a previously installed version of an update.

Perform any of the following functions on this page:

- Click Check Now to obtain the latest information from Palo Alto Networks.
- Click Install in the Action column to update to that version.
- Click **Revert** for a version to return to that version.

- Click **Release Notes** to view a description of an update.
- Click **Upload** to install a file that you previously stored on your PC. Browse to select the file, and click **Install from File**. Choose the file that you just selected from the drop-down list, and click **OK** to install.
- Click the **Schedule** link to schedule automatic updates. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select **Download Only**, you can install the downloaded update by clicking the **Install** link in the **Action** column on the **Dynamic Updates** page. When you click **OK**, the update is scheduled. No commit is required. You can also indicate how persistent the content must be (number of hours) for the action to take place and whether the upload should be synchronized to peer firewalls.

Administrator Roles, Profiles, and Accounts

The firewall supports the following options to authenticate administrative users who attempt to log in to the firewall:

- **Local database**—The user login and password information is entered directly into the firewall database.
- **RADIUS**—Existing Remote Authentication Dial In User Service (RADIUS) servers are used to authenticate users.
- LDAP—Existing Lightweight Directory Access Protocol (LDAP) servers are used to authenticate users.
- Kerberos—Existing Kerberos servers are used to authenticate users.
- Client Certificate—Existing client certificates are used to authenticate users.

When you create an administrative account, you specify local authentication or client certificate (no authentication profile), or an authentication profile (RADIUS, LDAP, Kerberos, or local DB authentication). This setting determines how the administrator's authentication is checked.

Administrator roles determine the functions that the administrator is permitted to perform after logging in. You can assign roles directly to an administrator account, or define role profiles, which specify detailed privileges, and assign those to administrator accounts.

Refer to the following sections for additional information:

- For instructions on setting up authentication profiles, refer to "Setting Up Authentication Profiles" on page 62.
- For instructions on setting up role profiles, refer to "Defining Administrator Roles" on page 58.
- For instructions on setting up administrator accounts, refer to "Creating Administrative Accounts" on page 59.
- For information on SSL virtual private networks (VPNs), refer to "Configuring GlobalProtect" on page 335.
- For instructions on defining virtual system domains for administrators, refer to "Specifying Access Domains for Administrators" on page 61.

• For instructions on defining certificate profiles for administrators, refer to "Certificate Profile" on page 89.

Username and Password Requirements

The following table lists the valid characters that can be used in usernames and passwords for PAN-OS and Panorama accounts.



Note: The following restrictions apply to PAN-OS 3.1 and later.

Account Type	Restrictions		
Password Character Set	There are no restrictions on any password field character sets.		
Remote Admin, SSL-VPN, or Captive Portal	The following characters are not allowed for the username: • Backtick (`) • Angular brackets (< and >) • Ampersand (&) • Asterisk (*) • At sign (@) • Question mark (?) • Pipe (1) • Single-Quote (') • Semicolon (;) • Double-Quote (") • Double-Quote (") • Dollar (\$) • Parentheses ('(' and ')') • Colon (':')		
Local Administrator Accounts	The following are the allowed characters for local usernames: • Lowercase (a-z) • Uppercase (A-Z) • Numeric (0-9) • Underscore (_) • Period (.) • Hyphen (-) <i>Note:</i> Login names cannot start with a hyphen (-).		

Table 7 Valid Characters for Usernames and Passwords



Note: The system does not impose restrictions on the response page character set.

Defining Administrator Roles

Device > Admin Roles

Use the **Admin Roles** page to define role profiles that determine the access and responsibilities available to administrative users. For instructions on adding administrator accounts, refer to "Creating Administrative Accounts" on page 59.

There are also three pre-defined Admin Roles that can be used for common criteria purposes. You first use the Superuser role for the initial configuration of the device and to create the administrator accounts for the Security Administrator, Audit Administrator, and Cryptographic Administrator. Once the accounts are created and the proper common criteria Admin Roles are applied, you then login using those accounts. The default Superuser account in FIPS or CC mode is **admin** and has a default password of **paloalto**. In standard operating mode, the default **admin** password is **admin**. The pre-defined Admin Roles were created where there is no overlap in capabilities, except that all have read-only access to the audit trail (except audit administrator with full read/delete access. These admin roles cannot be modified and are defined as follows:

- auditadmin—The Audit Administrator is responsible for the regular review of the firewall's audit data.
- cryptoadmin—The Cryptographic Administrator is responsible for the configuration and maintenance of cryptographic elements related to the establishment of secure connections to the firewall.
- securityadmin—The Security Administrator is responsible for all other administrative tasks (e.g. creating the firewall's security policy) not addressed by the other two administrative roles.

Field	Description		
Name	Enter a name to identify this administrator role (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.		
Description	Enter an optional description of the role (up to 255 characters).		
Role	Select the general scope of administrative responsibility from the drop- down list.		
WebUI	Click the icons for specified areas to indicate the type of access permitted		
	for the web interface:		
	• Read/write access to the indicated page.		
	 Read only access to the indicated page. 		
	• No access to the indicated page.		
CLI Role	Select the type of role for CLI access:		
	• disable—Access to the device CLI not permitted.		
	• superuser —Full access to the current device.		
	• superreader—Read-only access to the current device.		
	 deviceadmin—Full access to a selected device, except for defining new accounts or virtual systems. 		
	• devicereader —Read-only access to a selected device.		

Table 8. Administrator Role Settings

Defining Password Profiles

Device > Password Profiles

Password profiles allow you to set basic password requirements for an individual local account. If you enable **Minimum Password Complexity**, which provides password requirements for all local accounts, this password profile will override those settings. Refer to "Minimum Password Complexity" on page 35 for more information and refer to "Username and Password Requirements" on page 57 for information on valid characters that can be used for accounts.

To apply a password profile to an account, go to **Device > Administrators**, select an account and then choose the profile from the **Password Profile** drop-down.

Field	Description		
Name	Enter a name to identify the password profile (up to 31 characters). The nam is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.		
Required Password Change Period	Require that administrators change their password on a regular basis specified a by the number of days set, ranging from 0-365 days. Example, if the value is set to 90, administrators will be prompted to change their password every 90 days.		
(days)	You can also set an expiration warning from 0-30 days and specify a grace period.		
Expiration Warning Period (days)	If a required password change period is set, this setting can be used to prompt the user to change their password at each log in as the forced password change date approaches (range 0-30 days).		
Post Expiration Admin Login Count	Allow the administrator to log in the specified number of times after their account has expired. Example, if the value is set to 3 and their account has expired, they can log in 3 more times before their account is locked out (range 0-3 logins).		
Post Expiration Grace Period (days)	Allow the administrator to log in the specified number of days after their account has expired (range is 0-30 days).		

Table 9 Password Profile Settings

Creating Administrative Accounts

Device > Administrators

Administrator accounts control access to the firewall. Each administrator can have full or read-only access to a single device or to a virtual system on a single device. The predefined **admin** account has full access.

The following authentication options are supported:

- Password authentication—The user enters a user name and password to log in. No certificates are required.
- Client certificate authentication (web)—If you select this check box, a user name and password are not required; the certificate is sufficient to authenticate access to the firewall.

• Public key authentication (SSH)—The user can generate a public/private key pair on the machine that requires access to the firewall, and then upload the public key to the firewall to allow secure access without requiring the user enter a user name and password.



Note: To ensure that the device management interface remains secure, it is recommended that administrative passwords be changed periodically using a mixture of lower-case letters, upper-case letters, and numbers. You can also enforce "Minimum Password Complexity" from Setup > Management.

Field	Description	
Name	Enter a login name for the user (up to 15 characters). The name is case sensitive and must be unique. Use only letters, numbers, hyphens, periods, and underscores.	
	<i>Note:</i> Login names cannot start with a hyphen (-).	
Authentication Profile	Select an authentication profile for administrator authentication according to the settings in the specified authentication profile. This setting can be used for RADIUS, LDAP, Kerberos, or Local DB authentication.	
	For instructions on setting up authentication profiles, refer to "Setting Up Authentication Profiles" on page 62.	
Use only client certificate authentication (web)	Select the check box to use client certificate authentication for web access. If you select this check box, a user name and password are not required; the certificate is sufficient to authenticate access to the firewall.	
New Password Confirm New Password	Enter and confirm a case-sensitive password for the user (up to 15 characters). <i>You can also enforce "Minimum Password" from Setup > Management.</i>	
Use Public Key Authentication (SSH)	Select the check box to use SSH public key authentication. Click Import Key and browse to select the public key file. The uploaded key is displayed in the read-only text area.	
	Supported key file formats are IETF SECSH and OpenSSH. Supported key algorithms are DSA (1024 bits) and RSA (768-4096 bits).	
	Note: If the public key authentication fails, a username and password prompt is presented to the user.	

Table 10. Administrator Account Settings

Field	Description		
Role	Select an option for assigning a role to this user. The role determines what the user can view and modify.		
	If you choose Dynamic , you can select any of the following pre- specified roles from the drop-down list:		
	• Superuser—Full access to the current device.		
	• Superuser (read-only)—Read-only access to the current device.		
	 Device Admin—Full access to a selected device, except for defining new accounts or virtual systems. 		
	• Device administrator (read-only)—Read-only access to a selected device.		
	• Vsys Admin—Full access to a selected virtual system on a specific device (if multiple virtual systems are enabled).		
	• Vsys Admin (read-only)—Read-only access to a selected virtual system on a specific device.		
	 Role Based Admin—Access based on assigned roles, as defined in "Defining Administrator Roles" on page 58. 		
	If you choose Role Based , select a previously-defined role profile from the drop-down list. For instructions on defining role profiles, refer to "Defining Administrator Roles" on page 58.		
Virtual System	Select the virtual systems that you want the administrator to have access to, and click Add to move them from the Available area to the Selected area.		

Table 10. Administrator Account Settings (Continued)



Note: On the Panorama **Administrators** page for "superuser," a lock icon is shown in the right column if an account is locked out. The administrator can click the icon to unlock the account.

Specifying Access Domains for Administrators

Device > Access Domain

Use the **Access Domain** page to specify domains for administrator access to the firewall. The access domain is linked to RADIUS Vendor-specific Attributes (VSAs) and is supported only if a RADIUS server is used for administrator authentication. For information on configuring RADIUS, refer to "Configuring RADIUS Server Settings" on page 65.

When an administrator attempts to log in to the firewall, the firewall queries the RADIUS server for the administrator's access domain. If there is an associated domain on the RADIUS server, it is returned and the administrator is restricted to the defined virtual systems inside

the named access domain on the device. If RADIUS is not used, the access domain settings on this page are ignored. For information on Panorama Access Domains, refer to "Specifying Panorama Access Domains for Administrators" on page 392.

Field	Description
Name	Enter a name for the access domain (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, hyphens, and underscores.
Virtual Systems	Select virtual systems in the Available column and click Add to select them.
	<i>Note:</i> Access Domains are only supported on devices that support virtual systems.

Table 11	. A	ccess	Domain	Settings
----------	-----	-------	--------	----------

Authentication Profiles

Authentication profiles specify local database, RADIUS, LDAP, or Kerberos settings and can be assigned to administrator accounts, SSL-VPN access, and captive portal. When an administrator attempts to log in to the firewall directly or through an SSL-VPN or captive portal, the firewall checks the authentication profile that is assigned to the account and authenticates the user based on the authentication settings.

If the user does not have a local administrator account, the authentication profile that is specified on the device **Setup** page determines how the user is authenticated (refer to "Defining Management Settings" on page 30):

- If you specify **RADIUS** authentication settings on the **Setup** page and the user does not have a local account on the firewall, then the firewall requests authentication information for the user (including role) from the RADIUS server. The Palo Alto Networks RADIUS dictionary file containing the attributes for the various roles is available on the support site at *https://live.paloaltonetworks.com/docs/DOC-3189*.
- If **None** is specified as the authentication profile on the **Settings** page, then the user must be authenticated locally by the firewall according to the authentication profile that is specified for the user.

Setting Up Authentication Profiles

Device > Authentication Profile

Use the **Authentication Profile** page to configure authentication settings that can be applied to accounts to manage access to the firewall.

Field	Description
Name	Enter a name to identify the profile (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.

Table 12. Authentication Profile Settings

Field	Description
Lockout Time	Enter the number of minutes that a user is locked out if the number of failed attempts is reached (0-60 minutes, default 0). 0 means that the lockout is in effect until it is manually unlocked.
Failed Attempts	Enter the number of failed login attempts that are allowed before the account is locked out (1-10, default 0). 0 means that there is no limit.
Allow List	Specify the users and groups that are explicitly allowed to authenticate. Click Edit Allow List and do any of the following:
	• Select the check box next to the appropriate user or user group in the Available column, and click Add to add your selections to the Selected column.
	• Use the All check box to apply to all users.
	• Enter the first few characters of a name in the Search field to list all the users and user groups that start with those characters. Selecting an item in the list sets the check box in the Available column. Repeat this process as often as needed, and then click Add .
	• To remove users or user groups, select the appropriate check boxes in the Selected column and click Remove , or select any to clear all users.
Authentication	Choose the type of authentication:
	• None—Do not use any authentication on the firewall.
	• Local Database—Use the authentication database on the firewall.
	• RADIUS —Use a RADIUS server for authentication.
	• LDAP—Use LDAP as the authentication method.
	• Kerberos—Use Kerberos as the authentication method.
Server Profile	If you select RADIUS, LDAP, or Kerberos as the authentication method, choose the authentication server from the drop-down list. Servers are configured on the Server pages. Refer to "Configuring RADIUS Server Settings" on page 65, "Configuring LDAP Server Settings" on page 66, and "Configuring Kerberos Settings (Native Active Directory Authentication)" on page 67.
Login Attribute	If you selected LDAP as the authentication method, enter the LDAP directory attribute that uniquely identifies the user.

Table 12.	Authentication	Profile Settings	(Continued)
-----------	----------------	-------------------------	-------------

Field	Description
Password Expiry Warning	If you selected LDAP as the authentication method, enter the number of days prior to password expiration to send an automated message to the user. If the field is left blank, no warning is provided. This is supported for the following databases: Active Directory, eDirectory, and Sun ONE Directory.
	This setting is used for SSL-VPN. For more information, refer to "Configuring GlobalProtect" on page 335.
	You can customize the expiration warning message as part of the SSL-VPN login page by editing the script
	<script> function getPassWarnHTML(expdays) { var str = "Your password will expire in " + expdays + " days"; return str; } </script>
	Changing the value of the str variable changes the displayed message.

Table 12. Authentication Profile Settings (Continued)

Creating a Local User Database

You can set up a local database on the firewall to store authentication information for remote access users, device administrators, and captive portal users. No external authentication server is required with this configuration, so all account management is performed on the device or from Panorama.

When configuring Captive Portal, you first create the local account, add it to a User Group and create an Authentication Profile using the new group. You then enable Captive Portal from Device > User Authentication > Captive Portal and select the Authentication Profile. Once this is configured, you can then create a policy from Policies > Captive Portal. For more information, refer to "Captive Portals" on page 285.

Adding Local Users

Device > Local User Database > Users

Use the Local Users page to add user information to the local database.

Field	Description
Local User Name	Enter a name to identify the user (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Location	Choose a virtual system or choose Shared to make the certificate available to all virtual systems.
Mode	Use this field to specify the authentication option:
	• Password —Enter and confirm a password for the user.
	• Password Hash—Enter a hashed password string.
Enable	Select the check box to activate the user account.

Table 13. Local User Settings

Adding Local User Groups

Device > Local User Database > User Groups

Use the Local User Groups page to add user group information to the local database.

Table 14. Local User Group Settings

Field	Description
Local User Group Name	Enter a name to identify the group (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Location	Choose a virtual system or choose Shared to make the certificate available to all virtual systems.
All Local Users	Click Add to select the users you want to add to the group.

Configuring RADIUS Server Settings

► Device > Server Profiles > RADIUS

Use the **RADIUS** page to configure settings for the RADIUS servers that are identified in authentication profiles. Refer to "Authentication Profiles" on page 62.

Field	Description
Name	Enter a name to identify the server (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Location	Choose a virtual system, or choose Shared to make the profile available to all virtual systems.
Administrator Use Only	Use this server profile for administrator authentication only.
Domain	Enter the RADIUS server domain. The domain setting is used if the user does not specify a domain when logging in.
Timeout	Enter an interval after which an authentication request times out (1-30 seconds, default 3 seconds).
Retries	Enter the number of automatic retries following a timeout before the request fails (1-5, default 3).
Retrieve User Group	Select the check box to use RADIUS VSAs to define the group that has access to the firewall.
Servers	Configure information for each server in the preferred order.
	• Name—Enter a name to identify the server.
	• IP address—Enter the server IP address.
	• Port —Enter the server port for authentication requests.
	• Secret/Confirm Secret—Enter and confirm a key to verify and encrypt the connection between the firewall and the RADIUS server.

Table 15. RADIUS Server Settings

Configuring LDAP Server Settings

► Device > Server Profiles > LDAP

Use the **LDAP** page to configure settings for the LDAP servers to use for authentication by way of authentication profiles. Refer to "Authentication Profiles" on page 62.

Field	Description
Name	Enter a name to identify the profile (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Location	Choose a virtual system, or choose Shared to make the profile available to all virtual systems.
Administrator Use Only	Use this server profile for administrator authentication only.
Servers	Specify the host names, IPv4 or IPv6 addresses, and ports of your LDAP servers.
Domain	Enter the server domain name. This domain name should be the NetBIOS name of the domain and will be added to the username when authentication is performed. For example, if your domain is paloaltonetworks.com, you only need to enter paloaltonetworks.
Туре	Choose the server type from the drop-down list.
Base	Specify the root context in the directory server to narrow the search for user or group information.
Bind DN	Specify the login name (Distinguished Name) for the directory server.
Bind Password/ Confirm Bind Password	Specify the bind account password. The agent saves the encrypted password in the configuration file.
SSL	Select to use secure SSL or Transport Layer Security (TLS) communications between the Palo Alto Networks device and the directory server.
Time Limit	Specify the time limit imposed when performing directory searches (1 - 30 seconds, default 30 seconds).
Bind Time Limit	Specify the time limit imposed when connecting to the directory server (1 - 30 seconds, default 30 seconds).
Retry Interval	Specify the interval after which the system will try to connect to the LDAP server after a previous failed attempt (1-3600 seconds).

Table 16. LDAP Server Settings

Configuring Kerberos Settings (Native Active Directory Authentication)

Device > Server Profiles > Kerberos

Use the **Kerberos** page to configure Active Directory authentication without requiring customers to start Internet Authentication Service (IAS) for RADIUS support. Configuring a Kerberos server allows users to authenticate natively to a domain controller.

When the Kerberos settings are configured, Kerberos becomes available as an option when defining authentication profiles. Refer to "Authentication Profiles" on page 62.

You can configure the Kerberos settings to recognize a user account in any of the following formats, where domain and realm are specified as part of the Kerberos server configuration:

- domain\username
- username@realm
- username

Field	Description	
Name	Enter a name to identify the server (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.	
Location	Choose a virtual system, or choose Shared to make the profile available to all virtual systems.	
Administrator Use Only	Use this server profile for administrator authentication only.	
Realm	Specify the hostname portion of the user login name (up to 127 characters)	
	Example: The user account name <i>user@example.local</i> has realm <i>example.local</i> .	
Domain	Specify the domain for the user account (up to 63 characters).	
Servers	For each Kerberos server, click Add and specify the following settings:	
	• Server —Enter the server IP address.	
	• Host—Enter the server FQDN.	
	• Port —Enter an optional port number for communication with the server.	

Table 17. Kerberos Server Settings

Authentication Sequence

In some environments, user accounts reside in multiple directories. Guest or other accounts may also be stored in different directories. An authentication sequence is a set of authentication profiles that are applied in order when a user attempts to log in to the firewall. The firewall will always try the local database first, and then each profile in sequence until the user is identified. Access to the firewall is denied only if authentication fails for any of the profiles in the authentication sequence.

For example, after the local database is checked, you can configure an authentication sequence to then try RADIUS, followed by LDAP authentication.

Setting Up Authentication Sequences

Device > Authentication Sequence

Use the **Authentication Sequence** page to configure sets of authentication profiles that are tried in order when a user requests access to the firewall. The user is granted access if authentication is successful using any one of the authentication profiles in the sequence. For more information, refer to "Authentication Profiles" on page 62.

Field	Description
Profile Name	Enter a name to identify the profile (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems.
Lockout Time	Enter the number of minutes that a user is locked out if the number of failed attempts is reached (0-60 minutes, default 0). 0 means that the lockout is in effect until it is manually unlocked.
Failed Attempts	Enter the number of failed login attempts that are allowed before the account is locked out (1-10, default 0). 0 means that there is no limit.
Profile List	Choose the authentication profiles to include in the authentication sequence. To change the list order, select an entry and click Move Up or Move Down .

Table 18.	Authentication	Sequence	Settings
-----------	----------------	----------	----------

Firewall Logs

Monitor > Logs

The firewall provides logs that record configuration changes, system events, security threats, and traffic flows. For each log, you can enable remote logging to a Panorama server, and generate SNMP traps, syslog messages, and email notifications.

The following table describes the logs and logging options.

Log	Description
Alarms	The alarms log records detailed information on alarms that are generated by the system. The information in this log is also reported in the Alarms window. Refer to "Viewing Alarms" on page 85.
Configuration	The configuration log records each configuration change, including the date and time, the administrator user name, and whether the change succeeded or failed.
	All configuration log entries can be sent to Panorama, syslog, and email servers, but they cannot generate SNMP traps.
	To view expanded details about configuration log entries, move your cursor over the Before Change of After Change column and click the ellipsis symbol

Table 19 Log Types and Settings

Log	Description
Data Filtering	The data filtering log records information on the security policies that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall (refer to "Data Filtering Profiles" on page 223.
	If you configure a file blocking profile to block specific file types, the file type and file name will appear in the data filtering log, so you can see what was blocked.
HIP Match	The HIP match log lists the host information profile (HIP) match requests for GlobalProtect. Refer to "Configuring GlobalProtect" on page 335.
System	The system log records each system event, such as HA failures, link status changes, and administrators logging in and out. Each entry includes the date and time, the event severity, and an event description.
	System log entries can be logged remotely by severity level. For example, you can generate SNMP traps and email notifications for just critical and high-level events.
Threat	The threat log records each security alarm generated by the firewall. Each entry includes the date and time, the threat type, such as a virus or spyware/vulnerability filtering violation, the source and destination zones, addresses, and ports, the application name, and the action and severity.
	Threat log entries can be logged remotely by severity level by defining log forwarding profiles, and then assigning the profiles to security rules (refer to "Security Policies" on page 187). Threats are logged remotely only for the traffic that matches the security rules where the logging profile is assigned.
	Center (refer to "Reports and Logs" on page 251).
Traffic	The traffic log can record an entry for the start and end of each session. Each entry includes the date and time, the source and destination zones, addresses, and ports, the application name, the security rule applied to the session, the rule action (allow, deny, or drop), the ingress and egress interface, and the number of bytes.
	Each security rule specifies whether the start and/or end of each session is logged locally for traffic that matches the rule. The log forwarding profile assigned to the rule determines whether the locally logged entries are also logged remotely.
	Traffic logs are used in generating reports and in the Application Command Center (refer to "Reports and Logs" on page 251).
URL Filtering	The URL filtering log records entries for URL filters, which block access to specific web sites and web site categories or generate an alert when a user accesses a proscribed web site.
	If you are using the PAN-DB and wold like to request a URL category change for a website, open the URL Filtering log details for the desired log entry and click on Request Categorization Change . Select the suggested category from the drop-down menu, enter your email address (optional) and comments (optional) and click Send .
	If you are using the BrightCloud DB, click Request Categorization Change and you will be directed to the BrightCloud support page where you can submit you request.
	(refer to "URL Filtering Profiles" on page 217).

Table 19 Log Types and Settings (Continued)

Log	Description
WildFire	The WildFire logs will show the result of files analyzed by WildFire. Some of the fields that will be displayed include: Filename, Source Zone, Destination Zone, Attacker, Victim, and Application.
	Note: A WildFire subscription is required for integrated logging. If you do not have a subscription, you can view log information by using the WildFire portal at https://wildfire.paloaltonetworks.com.
	(refer to "About WildFire" on page 421 and "Using the WildFire Portal" on page 425)

Table 19 Log Types and Settings (Continued)

Logging Configuration

You can configure the firewall to send log entries to a Panorama centralized management system, SNMP trap sinks, syslog servers, and email addresses.

The following table describes the remote log destinations.

Destination	Description
Panorama	All log entries can be forwarded to a Panorama centralized management system. To specify the address of the Panorama server, refer to "Defining Management Settings" on page 30.
SNMP trap	SNMP traps can be generated by severity level for system, threat, and traffic log entries, but not for configuration log entries. To define the SNMP trap destinations, refer to "Configuring SNMP Trap Destinations" on page 75.
Syslog	Syslog messages can be generated by severity level for system, threat, and traffic log entries, and for all configuration log entries. To define the syslog destinations, refer to "Configuring Syslog Servers" on page 76.
Email	Emails can be generated by severity level for system, threat, and traffic log entries, and for all configuration log entries. To define the email addresses and servers, refer to "You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in Device > Server Profiles > Syslog. Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type." on page 77.

Table 20 Remote Log Destinations

Scheduling Log Exports

Device > Scheduled Log Export

You can schedule exports of logs and save them to a File Transfer Protocol (FTP) server in CSV format or use Secure Copy (SCP) to securely transfer data between the device and a remote host. Log profiles contain the schedule and FTP server information. For example, a profile may specify that the previous day's logs are collected each day at 3AM and stored on a particular FTP server.

When you click **OK** after creating a new entry, the new profile is added to the **Scheduled Log Export** page. You must commit the change for the export to take place. If you are using SCP, you need to click the **Test SCP server connection** button to test connectivity between the firewall and the SCP server and you must verify and accept the host key of the SCP server.

Field	Description
Name	Enter a name to identify the profile (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
	You cannot change the name after the profile is created.
Description	Enter an optional description (up to 255 characters).
Enabled	Select the check box to enable the scheduling of log exports.
Log Type	Select the type of log (traffic, threat, url, data, or hipmatch). Default is traffic.
Scheduled export start time (daily)	Enter the time of day (hh:mm) to start the export, using a 24-hour clock (00:00 - 23:59).
Protocol	Select the protocol to use to export logs from the firewall to a remote host. You can use SCP to export logs securely, or you can use FTP, which is not a secure protocol.
Hostname	Enter the host name or IP address of the FTP server that will be used for the export.
Port	Enter the port number that the FTP server will use. Default is 21.
Path	Specify the path located on the FTP server that will be used to store the exported information.
Enable FTP Passive Mode	Select the check box to use passive mode for the export. By default, this option is selected.
Username	Enter the user name for access to the FTP server. Default is anonymous.
Password	Enter the password for access to the FTP server. A password is not required if the user is "anonymous."

Table 21. Scheduled Log Export Settings

_

Defining Configuration Log Settings

Device > Log Settings > Config

The configuration log settings specify the configuration log entries that are logged remotely with Panorama, and sent as syslog messages and/or email notifications.

Field	Description
Panorama	Select the check box to enable sending configuration log entries to the Panorama centralized management system.
SNMP Trap	To generate SNMP traps for configuration log entries, select trap name. To specify new SNMP trap destinations, refer to "Configuring SNMP Trap Destinations" on page 75.
Email	To generate email notifications for configuration log entries, select an email profile from the drop-down menu. To specify a new email profile, refer to "Configuring Email Notification Settings" on page 83.
Syslog	To generate syslog messages for configuration log entries, select the name of the syslog server. To specify new syslog servers, refer to "Configuring Syslog Servers" on page 76.

Table 22. Configuration Log Settings

Defining System Log Settings

Device > Log Settings > System

The system log settings specify the severity levels of the system log entries that are logged remotely with Panorama and sent as SNMP traps, syslog messages, and/or email notifications. The system logs show system events such as HA failures, link status changes, and administrators logging in and out.

Table 23.System Log Settings

Field	Description
Panorama	Select the check box for each severity level of the system log entries to be sent to the Panorama centralized management system. To specify the Panorama server address, refer to "Defining Management Settings" on page 30.
	The severity levels are:
	• Critical—Hardware failures, including HA failover, and link failures.
	 High—Serious issues, including dropped connections with external devices, such as syslog and RADIUS servers.
	• Medium—Mid-level notifications, such as antivirus package upgrades.
	• Low—Minor severity notifications, such as user password changes.
	• Informational —Log in/log off, administrator name or password change, any configuration change, and all other events not covered by the other severity levels.
Field	Description
------------------------------	---
SNMP Trap Email Syslog	Under each severity level, select the SNMP, syslog, and/or email settings that specify additional destinations where the system log entries are sent. To define new destinations, refer to:
	• "Configuring SNMP Trap Destinations" on page 75.
	• "Configuring Syslog Servers" on page 76.
	• Configuring Email Noulication Settings on page 83.

Table 23. System Log Settings (Continued)

Defining HIP Match Log Settings

Device > Log Settings > HIP Match

The Host Information Profile (HIP) match log settings are used to provide information on security policies that apply to GlobalProtect clients. For more information, refer to "Overview" on page 335.

Field	Description
Panorama	Select the check box to enable sending configuration log entries to the Panorama centralized management system.
SNMP Trap	To generate SNMP traps for HIP match log entries, select the name of the trap destination. To specify new SNMP trap destinations, refer to "Configuring SNMP Trap Destinations" on page 75.
Email	To generate email notifications for configuration log entries, select the name of the email settings that specify the appropriate email addresses. To specify new email settings, refer to "You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in Device > Server Profiles > Syslog. Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type." on page 77.
Syslog	To generate syslog messages for configuration log entries, select the name of the syslog server. To specify new syslog servers, refer to "Configuring Syslog Servers" on page 76.

Table 24. HIP Match Log Settings

Defining Alarm Log Settings

Device > Log Settings > Alarms

Use the **Alarms** page to configure notifications when a security rule (or group of rules) has been hit repeatedly in a set period of time.

Field	Description
Enable Alarms	Enable alarms based on the events listed on this page.
Enable CLI Alarm Notifications	Enable CLI alarm notifications whenever alarms occur.

Table 25. Alarm Log Settings

Field	Description
Enable Web Alarm Notifications	Open a window to display alarms on user sessions, including when they occur and when they are acknowledged.
Enable Audible Alarms	The firewall will continuously play an audible tone when unacknowledged alarms exist in the web interface or CLI.
Encryption/Decryption Failure Threshold	Specify the number of encryption/decryption failures after which an alarm is generated.
Log DB Alarm Threshold (% Full)	Generate an alarm when a log database reaches the indicated percentage of the maximum size.
Security Policy Limits	An alarm is generated if a particular IP address or port hits a deny rule the number of times specified in the Security Violations Threshold setting within the period (seconds) specified in the Security Violations Time Period setting.
Security Policy Group Limits	An alarm is generated if the collection of rules reaches the number of rule limit violations specified in the Violations Threshold field during the period specified in the Violations Time Period field. Violations are counted when a session matches an explicit deny policy.
	Use Security Policy Tags to specify the tags for which the rule limit thresholds will generate alarms. These tags become available to be specified when defining security policies.
Selective Audit	<i>Note:</i> These settings appear on the <i>Alarms</i> page only in Common Criteria mode.
	Specify the following settings:
	• CC Specific Logging —Enables verbose logging required for Common Criteria (CC) compliance.
	• Login Success Logging—Logs the success of administrator logins to the firewall.
	• Login Failure Logging—Logs the failure of administrator logins to the firewall.
	• Suppressed Administrators —Does not generate logs for changes that the listed administrators make to the firewall configuration.

Table 25. Alarm Log Settings (Continued)

Managing Log Settings

Device > Log Settings > Manage Logs

Click the links on this page to clear the indicated logs.

Configuring SNMP Trap Destinations

Device > Server Profiles > SNMP Trap

To generate SNMP traps for system, traffic, or threat logs, you must specify one or more SNMP trap destinations. After you define the trap destinations, you can use them for system log entries (refer to "Defining System Log Settings" on page 72).

Field	Description
Name	Enter a name for the SNMP profile (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems.
Version	Choose the SNMP version (V2c or V3).
V2c settings	If you choose V2c, configure the following settings:
	• Server —Specify a name for the SNMP trap destination name (up to 31 characters).
	• Manager—Specify the IP address of the trap destination.
	• Community —Specify the community string required to send traps to the specified destination (default public).
V3 settings	If you choose V3, configure the following settings:
	• Server—Specify the SNMP trap destination name (up to 31 characters).
	• Manager—Specify the IP address of the trap destination.
	• User—Specify the SNMP user.
	 EngineID—Specify the engine ID of the firewall. The input is a string in hexadecimal representation. The engine ID is any number between 5 to 64 bytes. When represented as a hexadecimal string this is between 10 to 128 characters (2 characters for each byte) with two additional characters for 0x that you need to use as a prefix in the input string. Each firewall has a unique engine ID, which you can get by using a MIB browser to run a GET for OID 1.3.6.1.6.3.10.2.1.1.0.
	• Auth Password—Specify the user's authentication password (min- imum 8 characters, maximum of 256 characters, and no character restrictions). All characters allowed). Only Secure Hash Algorithm (SHA) is supported.
	• Priv Password —Specify the user's encryption password (minimum 8 characters, maximum of 256 characters, and no character restrictions). Only Advanced Encryption Standard (AES) is supported.

Table 26. SNMP Trap Destination Settings



Note: Do not delete a destination that is used in any system log settings or logging profile.

SNMP MIBs

The firewall supports the following SNMP MIBs:

- "RFC 1213: MIB-II Support for The System Group, The Interfaces Group.
- "RFC 2863: IF-MIB The Interfaces Group MIB
- "RFC 2790: HOST-RESOURCES-MIB Support for hrDeviceTable and hrProcessorTable.
- "RFC 3433: ENTITY-SENSOR-MIB Support for entPhySensorTable.
- PAN-PRODUCT-MIB
- PAN-COMMON-MIB
- PAN-TRAPS-MIB
- PAN-LC-MIB

The full set of Enterprise MIBs is available in the Technical Documentation section on the Palo Alto Networks site at *https://live.paloaltonetworks.com/community/documentation*.

Configuring Syslog Servers

Device > Server Profiles > Syslog

To generate syslog messages for system, configuration, traffic, threat, or HIP match logs, you must specify one or more syslog servers. After you define the syslog servers, you can use them for system and configuration log entries (refer to "Defining System Log Settings" on page 72).

Field	Description
Name	Enter a name for the syslog profile (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems.
Servers Tab	
Name	Click Add and enter a name for the syslog server (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Server	Enter the IP address of the syslog server.
Port	Enter the port number of the syslog server (the standard port is 514).
Facility	Choose a level from the drop-down list.

Table 27. New Syslog Server

Field	Description
Custom Log Format Tab	
Log Type	Click the log type to open a dialog box that allows you to specify a custom log format. In the dialog box, click a field to add it to the Log Format area. Other text strings can be edited directly in the Log Format area. Click OK to save the settings.
	For details on the fields that can be used for custom logs, refer to "Custom Syslog Field Descriptions" on page 77.
Escaping	Specify escape sequences. Use the Escaped characters box to list all the characters to be escaped without spaces.

Table 27. New Syslog Server (Continued)



Note: You cannot delete a server that is used in any system or configuration log settings or logging profiles.

Custom Syslog Field Descriptions

You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in **Device > Server Profiles > Syslog.** Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type.

Field	Meaning
actionflags	A bit field indicating if the log was forwarded to Panorama. Available in PAN-OS 4.0.0 and above.
admin	User name of the Administrator performing the configuration.
after-change-detail	Details of the configuration after a change is made.
before-change-detail	Details of the configuration before a change is made.
formtted-receive_time	Time the log was received at the management plane, shown in CEF compliant time format.
cef-formatted-time_generated	Time the log was generated, shown in CEF compliant time format.
client	Client used by the Admin; Values are Web and CLI.
cmd	Command performed by the Admin; Values are add, clone, commit, delete, edit, move, rename, set, validate.
host	Host name or IP address of the client machine
path	The path of the configuration command issued. Up to 512 bytes in length.
receive_time	Time the log was received at the management plane.
result	Result of the configuration action. Values are Submitted, Succeeded, Failed, and Unauthorized.

Table 28 Config Fields

Field	Meaning
seqno	A 64bit log entry identifier incriminated sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above.
serial	Serial number of the device that generated the log
subtype	Subtype of the Config log; Unused.
time_generated	Time the log was received on the data plane.
type	Specifies type of log; Values are traffic, threat, config, system and hip-match.
vsys	Virtual System associated with the configuration log.

Table 28 Config Fields

Field	Meaning
actionflags	A bit field indicating if the log was forwarded to Panorama. Available in PAN-OS 4.0.0 and above.
cef-formatted-receive_time	Time the log was received at the management plane, shown in CEF compliant time format.
cef-formatted-time_generated	Time the log was generated, shown in CEF compliant time format.
eventid	String showing the name of the event
fmt	Detailed description of the event. Length is up to 512 bytes.
module	This field is valid only when the value of the Subtype field is general; It provides additional information about the sub-system generating the log. Values are general, management, auth, ha, upgrade, chassis.
number-of-severity	Severity level as an integer - informational-1, low-2, medium-3, high-4, critical-5.
object	Name of the object associated with the system log.
opaque	Detailed description of the event. Length is up to 512 bytes.
receive_time	Time the log was received at the management plane
seqno	A 64bit log entry identifier that increments sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above.
serial	Serial number of the device that generated the log
severity	Severity associated with the event; Values are informational, low, medium, high, critical
subtype	Subtype of the system log. Refers to the system daemon generating the log; Values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
time_generated	Time the log was received on the data plane.

Table 29 System Fields

Field	Meaning
type	Specifies type of log; Values are traffic, threat, config, system and hip-match.
vsys	Virtual System associated with the system event

Table 29 System Fields

Table 30 Threat Fields

Field	Meaning
action	Action taken for the session; Values are alert, allow, deny, drop, drop-all-packets, reset-client, reset-server, reset-both, block-url. See Action Field table below for meaning of each value.
actionflags	A bit field indicating if the log was forwarded to Panorama. Available in PAN-OS 4.0.0 and above.
app	Application associated with the session.
category	For URL Subtype, it is the URL category; for WildFire subtype, it is the verdict on the file and is either "malicious" or "benign"; for other subtypes, the value is "any".
cef-formatted-receive_time	Time the log was received at the management plane, shown in CEF compliant time format.
cef-formatted-time_generated	Time the log was generated, shown in CEF compliant time format.
contenttype	Content type of the HTTP response data. Maximum length 32 bytes. Applicable only when Subtype is URL. Available in PAN-OS 4.0.0 and above.
direction	Indicates the direction of the attack, 'client-to-server' or 'server- to-client'.
dport	Destination port utilized by the session.
dst	Original session destination IP address.
dstloc	Destination country or Internal region for private addresses. Maximum length is 32 bytes. Available in PAN-OS 4.0.0 and above.
dstuser	User name of the user to which the session was destined.
flags	32 bit field that provides details on the session; See Flags Field table for meaning of each value.
from	Zone the session was sourced from.
inbound_if	Interface that the session was sourced form.
logset	Log Forwarding Profile that was applied to the session.
misc	The actual URI when the subtype is URL; File name or file type when the subtype is file; and File name when the subtype is virus; File name when the subtype is wildfire. Length is 63 characters in PAN-OS versions before 4.0. From version 4.0, it is variable length with a maximum of 1023 characters.
natdport	Post-NAT destination port.

Field	Meaning	
natdst	If Destination NAT performed, the post-NAT Destination IP address.	
natsport	Post-NAT source port.	
natsrc	If Source NAT performed, the post-NAT Source IP address.	
number-of-severity	Severity level as an integer - informational-1, low-2, medium-3, high-4, critical-5.	
outbound_if	Interface that the session was destined to.	
proto	IP protocol associated with the session.	
receive_time	Time the log was received at the management plane.	
repeatcnt	Number of logs with same Source IP, Destination IP, and Threat ID seen within 5 seconds; Applies to all Subtypes except URL.	
rule	Name of the rule that the session matched.	
seqno	A 64bit log entry identifier that increments sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above.	
serial	Serial number of the device that generated the log.	
sessionid	An internal numerical identifier applied to each session.	
severity	Severity associated with the threat; Values are informational, low, medium, high, critical.	
sport	Source port utilized by the session.	
src	Original session source IP address.	
srcloc	Source country or Internal region for private addresses. Maximum length is 32 bytes. Available in PAN-OS 4.0.0 and above.	
srcuser	User name of the user that initiated the session.	
subtype	Subtype of threat log; Values are URL, virus, spyware, vulnerability, file, scan, flood, data, and wildfire.	
threatid	Palo Alto Networks identifier for the threat. It is a description string followed by a numerical identifier in parenthesis for some Subtypes. The numerical identifier is a 64-bit number from PAN- OS 5.0 and later.	
time_generated	Time the log was generated on the data plane.	
time_received	Time the log was received on the data plane.	
to	Zone the session was destined to.	
type	Specifies type of log; Values are traffic, threat, config, system and hip-match.	
vsys	Virtual System associated with the session.	
wildfire	Logs generated by WildFire.	

Field	Meaning	
action	Action taken for the session; Values are allow or deny. See Action Field table.	
actionflags	A bit field indicating if the log was forwarded to Panorama. Available from PAN-OS 4.0.0.	
app	Application associated with the session.	
bytes	Number of total bytes (transmit and receive) for the session.	
bytes_received	Number of bytes in the server-to-client direction of the session. Available from PAN-OS 4.1.0 on all models except the PA-4000 series.	
bytes_sent	Number of bytes in the client-to-server direction of the session. Available from PAN-OS 4.1.0 on all models except the PA-4000 series.	
category	URL category associated with the session (if applicable).	
cef-formatted-receive_time	Time the log was received at the management plane, shown in CEF compliant time format.	
cef-formatted-time_generated	Time the log was generated, shown in CEF compliant time format.	
dport	Destination port utilized by the session.	
dst	Original session destination IP address.	
dstloc	Destination country or Internal region for private addresses. Maximum length is 32 bytes. Available in PAN-OS 4.0.0 and above.	
dstuser	User name of the user to which the session was destined.	
elapsed	Elapsed time of the session.	
flags	32 bit field that provides details on session; See Flags Field table for meaning of each value. This field can be decoded by AND-ing the values with the logged value.	
from	Zone the session was sourced from.	
inbound_if	Interface that the session was sourced form.	
logset	Log Forwarding Profile that was applied to the session.	
natdport	Post-NAT destination port.	
natdst	If Destination NAT performed, the post-NAT Destination IP address.	
natsport	Post-NAT source port.	
natsrc	If Source NAT performed, the post-NAT Source IP address.	
outbound_if	Interface that the session was destined to.	
packets	Number of total packets (transmit and receive) for the session.	
pkts_received	Number of server-to-client packets for the session. Available from PAN-OS 4.1.0 on all models except the PA-4000 series.	
pkts_sent	Number of client-to-server packets for the session. Available from PAN-OS 4.1.0 on all models except the PA-4000 series.	

Table 31 Traffic Fields

Field	Meaning	
proto	IP protocol associated with the session.	
receive_time	Time the log was received at the management plane.	
repeatcnt	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; Used for ICMP only.	
rule	Name of the rule that the session matched.	
seqno	A 64bit log entry identifier incremented sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above.	
serial	Serial number of the device that generated the log.	
sessionid	An internal numerical identifier applied to each session.	
sport	Source port utilized by the session.	
src	Original session source IP address.	
srcloc	Source country or Internal region for private addresses. Maximum length is 32 bytes. Available in PAN-OS 4.0.0 and above.	
srcuser	User name of the user that initiated the session.	
start	Time of session start.	
subtype	Subtype of traffic log; Values are start, end, drop, and deny. See Subtype Field table for meaning of each value.	
time_generated	Time the log was generated on the data plane.	
time_received	Time the log was received on the data plane.	
to	Zone the session was destined to.	
type	Specifies type of log; Values are traffic, threat, config, system and hip-match.	
vsys	Virtual System associated with the session.	

Table 32 HIP Match Fields

Field	Meaning	
actionflags	A bit field indicating if the log was forwarded to Panorama. Available in PAN-OS 4.0.0 and above.	
cef-formatted-receive_time	Time the log was received at the management plane, shown in CEF compliant time format.	
cef-formatted-time_generated	Time the log was generated, shown in CEF compliant time format.	
machinename	Name of the Users machine.	
matchname	Name of the HIP Object or Profile.	
matchtype	Specifies whether the HIP field represents a HIP Object or a HIP Profile.	
receive_time	Time the log was received at the management plane.	

Field	Meaning	
repeatcnt	Number of times the HIP profile matched.	
seqno	A 64bit log entry identifier incremented sequentially. Each log type has a unique number space. Available in PAN-OS 4.0.0 and above.	
serial	Serial number of the device that generated the log.	
src	IP address of the source user.	
srcuser	User name of the Source user.	
subtype	Subtype of hip-match log; Unused.	
time_generated	Time the log was generated on the data plane.	
type	Specifies type of log; Values are traffic, threat, config, system and hip-match.	
vsys	Virtual System associated with the HIP Match log.	

Table 32 HIP Match Fields

Configuring Email Notification Settings

Device > Server Profiles > Email

To generate email messages for logs, you must configure an email profile. After you define the email settings, you can enable email notification for system and configuration log entries (refer to "Defining System Log Settings" on page 72). For information on scheduling email report delivery, refer to "Scheduling Reports for Email Delivery" on page 273.

Table 33.	Email	Notification	Settings
-----------	-------	--------------	----------

Field	Description	
Name	Enter a name for the email settings (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.	
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.	
Servers Tab		
Server	Enter a name to identify the server (1-31 characters). This field is just a label and does not have to be the host name of an existing SMTP server.	
Display Name	Enter the name shown in the From field of the email.	
From	Enter the From email address, such as "security_alert@company.com".	
То	Enter the email address of the recipient.	
Additional Recipient	Optionally, enter the email address of another recipient. You can only add one additional recipient. To add multiple recipients, add the email address for a distribution list.	
Gateway	Enter the IP address or host name of the Simple Mail Transport Protocol (SMTP) server used to send the email.	

Field	Description	
Custom Log Format Tab		
Log Type	Click the log type to open a dialog box that allows you to specify a custom log format. In the dialog box, click a field to add it to the Log Format area. Click OK to save the settings.	
Escaping	Include escaped characters and specify the escape character or characters.	

Table 33. Email Notification Settings (Continued)



Note: You cannot delete an email setting that is used in any system or configuration log settings or logging profiles.

Viewing Alarms

You can view the current list of alarms at any time by clicking the **Alarms** icon **Alarms** in the lower right corner of the web interface when the Alarm option is configured. This opens a window that lists the unacknowledged and acknowledged alarms in the current alarms log. To acknowledge alarms, select their check boxes and click **Acknowledge**. This action moves the alarms to the Acknowledged Alarms list. The alarms window also includes paging, column sort, and refresh controls.

The Alarms button is visible only when the **Enable Alarms** check box is selected on the **Device > Log Settings > Alarms > Alarm Settings** page.

Configuring Netflow Settings

Device > Server Profiles > Netflow

The firewall can generate and export NetFlow Version 9 records with unidirectional IP traffic flow information to an outside collector. NetFlow export can be enabled on any ingress interface in the system. Separate template records are defined for IPv4, IPv4 with NAT, and IPv6 traffic, and PAN-OS specific fields for App-ID and User-ID can be optionally exported. This feature is available on all platforms, except for the PA-4000 Series models.

The firewall supports the standard NetFlow templates and selects the correct one based on the data to be exported.

To configure NetFlow data exports, define a NetFlow server profile, which specifies the frequency of the export along with the NetFlow servers that will receive the exported data.

Then when you assign the profile to an existing firewall interface, all traffic flowing over that interface is exported to the specified servers. All interface types support assignment of a NetFlow profile. Refer to "Firewall Interfaces" on page 127 for information on assigning a NetFlow profile to an interface.

Field	Description
Name	Enter a name for the Netflow settings (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Template Refresh Rate	Specify the number of minutes or number of packets after which the Netflow template is refreshed (minutes range 1-3600, default 30 min; packets range 1-600, default 20).
Active Timeout	Specify the frequency at which data records are exported for each session (minutes).
Export PAN-OS Specific Field Types	Export PAN-OS specific fields such as App-ID and User-ID in Netflow records.
Servers	
Name	Specify a name to identify the server (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.

Table 34.	Netflow	Settings
-----------	---------	----------

Field	Description
Server	Specify the host name or IP address of the server. You can add a maximum of two servers per profile.
Port	Specify the port number for server access (default 2055).

Table 34. Netflow Settings (Continued)

Importing, Exporting and Generating Security Certificates

Certificates

Device > Certificate Management > Certificates

The **Certificates** page allows you to generate the following security certificates:

- Forward Trust—This certificate is presented to clients during decryption when the server to which they are connecting is signed by a CA in the firewall's trusted CA list. If a self-signed certificate is used for forward proxy decryption, you must click the certificate name in the Certificates page and select the Forward Trust Certificate check box.
- **Forward Untrust**—This certificate is presented to clients during decryption when the server to which they are connecting is signed by a CA that is not in the firewall's trusted CA list.
- **Trusted Root CA**—The certificate is marked as a trusted CA for forward decryption purposes.

When the firewall decrypts traffic, it checks the upstream certificate to see if it is issued by a trusted CA. If not, it uses a special untrusted CA certificate to sign the decryption certificate. In this case, the user sees the usual certificate error page when accessing the firewall and must dismiss the login warning.

The firewall has a large list of existing trusted CAs. The trusted root CA certificate is for additional CAs that are trusted for your enterprise but are not part of the pre-installed trusted list.

- **SSL Exclude**—This certificate excludes connections if they are encountered during SSL forward proxy decryption.
- **Certificate for Secure Web GUI**—This certificate authenticates users for access to the firewall web interface. If this check box is selected for a certificate, the firewall will use this certificate for all future web-based management sessions following the next commit operation.

Perform any of the following functions on the **Certificates** page:

- To revoke a certificate:
 - a. Select the certificate that you want to revoke.

b. Click **Revoke** and the certificate will be instantly set to the revoked status. No commit is required.

You can also click an existing certificate and click the Revoke icon.

- To renew a certificate:
 - a. Select the certificate that you want to renew.
 - b. Click **Renew** and then set the number of days to extend the certificate and click **OK**. A new certificate will be generated with the same attributes, but with a new serial number. The old certificate is then replaced with the new certificate.
- To import a web interface, trusted CA, or SSL forward proxy certificate:
 - a. Click Import.
 - b. Enter a name to identify the certificate.
 - c. Select the certificate file. If importing a PKCS #12 certificate and private key, this will be the single file holding both objects. If using PEM, this will be the public certificate only.
 - d. Click the **Import Private Key** check box to load the private key and enter the passphrase twice. If using the PKCS #12, the key file was selected above. If using PEM, browse to the encrypted private key file (generally named *.key).
 - e. Select the virtual system to which you want to import the certificate from the dropdown list.
- To export a certificate:
 - a. Select the certificate you want to export.
 - b. Click Export.
 - c. Choose the file format you would like the exported certificate to use (.pfx for PKCS#12 or .pem).
 - d. Select the **Export Private Key** check box and enter a passphrase twice to export the private key in addition to the certificate.
- Click **Save** and choose a location to copy the file to your local computer. To generate a certificate:
 - a. Click **Generate** to open the Generate Certificate window and specify the information described in the following table.
 - b. After generating the certificate, click the certificate link and specify the certificate type (Forward Trust, Forward Untrust, Trusted Root CA, SSL Exclude, or Certificate for Secure Web GUI).



Note: If you are using Panorama, you also have the option of generating a selfsigned certificate for the Panorama server. Refer to "Central Device Management Using Panorama" on page 381 for information on Panorama. • To import keys for high availability (HA), click **Import HA Key** and browse to specify the key file for import. To export keys for HA, click **Export HA Key** and specify a location to save the file. The HA keys must be swapped across the two firewalls. In other words, the Key from firewall 1 must be exported and then imported to firewall 2 and vice versa.

Field	Description		
Certificate Name	Enter a name (up to 31 characters) to identify the certificate. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required.		
Common Name	Enter the IP address or FQDN that will appear on the certificate.		
Location	Choose a virtual system or choose Shared to make the certificate available to all virtual systems.		
Signed By	Choose from a list of the CA certificates that were generated on the firewall. The selected certificate can be used to sign the certificate that is being created.		
Certificate Authority	Mark this certificate as a CA so that it can be used to sign other certificates on the firewall.		
OCSP Responder	Select an OSCP responder profile from the drop-down list. The profile is configured in Device > Certificate Management > OCSP Responder. When generating a certificate and entering an OCSP Responder, a look up will be performed for the host name of the IP address to generate a OCSP Responder URL, which will then appear in this drop-down.		
Number of Bits	Choose the key length for the certificate.		
Digest	Choose the digest algorithm for the certificate.		
Expiration (days)	Specify the number of days that the certificate will be valid. The default is 365 days.		
	If you specify a Validity Period in a GlobalProtect Portal Satellite configuration, that value will override the value entered in this field.		
Country State Locality Organization Department Email	Optionally specify additional information to identify the certificate. To view a list of country code definitions, click the ISO 6366 Country Codes link.		

Table 35. Settings to Generate a Certificate

Default Trusted Certificate Authorities

Device > Certificate Management > Certificates > Default Trusted Certificate Authorities

Use this page to control the certificate authorities (CAs) that the firewall will trust. You can disable and enable CAs as needed.

Field	Description	
EnableIf you have disabled a CA and want to enable it, click box next to the CA and then click Enable .		
Disable	Click the check box next to the CA that you want to disable, then click Disable . This may be desired if you only want to trust certain CAs, or remove all of them to only trust your local CA.	
Export	Click the check box next to the CA, then click Export to export the CA certificate. You can do this to import into another system, or if you want to view the certificate offline.	

Table 36 Trusted Certificate Authorities Settings

Certificate Profile

► Device > Certificate Management > Certificate Profile

You can create certificate profiles and then attach a profile to an administrator login on the **Setup** page or to an SSL-VPN login for use in authentication or with captive portals. You can also create certificate profiles to be used by GlobalProtect Gateways for authentication. Refer to "Defining Management Settings" on page 30 and "Captive Portals" on page 285.

Page Type	Description
Name	Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Location	If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems.
Username Field	Choose a user name option from the drop-down list.
Domain	Enter the domain for the profile.
CA Certificates	Choose a CA certificate from the drop-down list, specify the default OCSP URL, select an option to verify the CA certificate, and click Add . Repeat to add additional certificates. The OCSP Verify CA Certificate drop down is used if you have a separate OCSP responder that can verify certificates
Use CRL	Select the check box to use a Certificate Revocation List (CRL).
Use OCSP	Select the check box to use Online Certificate Status Protocol (OCSP) server. OCSP takes precedence over CRL.
CRL Receive Timeout	Specify an interval after which CRL requests time out (1 - 60 secs).
OCSP Receive Timeout	Specify an interval after which OCSP requests time out (1 - 60 secs).
Certificate Status Timeout	Specify an interval after which requests for certificate status time out (1 - 60 secs).

Table 37. Certificate Profile Settings

Page Type	Description
Block session if certificate status is unknown	Select the check box to block a sessions if the certificate status is unknown.
Block sessions if certificate status cannot be retrieved within timeout	Select the check box to block a session if the certificate status cannot be retrieved within the timeout interval.

Table 37. Certificate Profile Settings (Continued)

OCSP Responder

▶ Device > Certificate Management > OCSP Responder

Use the **OCSP Responder** (Online Certificate Status Protocol Responder) page to define a server that will be used to verify the revocation status of certificates issues by the PAN-OS device. When generating new certificates, you can specify the OCSP Responder that will be used.

To enable OCSP, go to **Device > Setup > Sessions** and under **Sessions Features** click **Decryption Certificate Revocation Settings**.

Field	Description
Name	Enter a name to identify the OCSP responder server (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Host Name	Enter the host name of the OCSP responder server that will be used to check certificate revocation status for your devices.

Table 38 OCSP Responder Settings

Encrypting Private Keys and Passwords on the Firewall

Device > Master Key and Diagnostics

Use the **Master Key and Diagnostics** page to specify a master key to encrypt private keys on the firewall. Private keys are stored in encrypted form by default even if a new master key is not specified. If the firewall is in Common Criteria mode, several crypto diagnostics capabilities are available. These diagnostics allow you to run scheduled crypto self-tests and on-demand self-tests.

To create a new master key, enter a 16 character string in the **New Master Key** field and then confirm the key. Enter a lifetime and reminder value and then click **OK**. You will need to update the master key before the expiration. For information on updating master keys, refer to "Updating Master Keys" on page 91.

- "Master Key and Diagnostic Settings" on page 91
- "Updating Master Keys" on page 91

Master Key and Diagnostic Settings

Field	Description
Current Master Key	Specify the key that is currently used to encrypt all of the private keys and passwords on the firewall.
New Master Key	To change the master key, enter and confirm a new key.
Confirm Master Key	
Life Time	Specify the number of days and hours after which the master key expires (range 1-730 days).
Time for Reminder	Specify the number of days and hours before expiration when the user is notified of the impending expiration (range 1-365 days).
Common Criteria	In Common Criteria mode, additional buttons are available to run a cryptographic algorithm self-test and software integrity self-test. A scheduler is also included to specify the times at which the two self-tests will run.

Table 39. Master Key and Diagnostics Settings

Updating Master Keys

When creating a new master keys an expiration period will be defined, so it's important to update the keys before the expiration period is reached. When using a master key in a high availability (HA) configuration, it is also important that the same master key is used on both devices to ensure that private keys and certificates are encrypted with the same master key. If the master keys are different, HA configuration synchronization will not work properly. The following section describes how to update the master key if they working properly and instructions are provided if they have already expired or if the master keys on each device in the HA pair are out of sync.

As of PAN-OS 5.0, if the master keys in an HA pair do not match, a critical system log will be generated.

Reasons for updating the master key:

- You want to change the default master key or change a master key that you created.
- Master keys in an HA configuration are out of sync.
- Master key is going to expire soon.

HA Master key update (keys are in sync and have not reached the expiration date):

- If you are updating the master keys and they are not expired or out of sync, before you
 update the keys, you need to commit the configuration and ensure that no pending
 configuration updates exist on both devices in the HA pair.
 You can view the commit status by clicking the Tasks link on the bottom right of the web
 interface on each device. From the CLI, run show jobs all to see all running jobs or
 show jobs pending to view pending jobs. To check that no pending updates are
 present, from configure mode, run check pending-changes and you should see No.
- 2. Disable configuration synchronization by navigating to **Device > High Availability** on both devices and from the **General** tab deselect the **Enable Config Sync** check box. Commit the configuration on both devices.
- 3. Update the master key on device A using 16 characters and commit the configuration.

- 4. Update the master key on device B with the same master key and commit the configuration.
- 5. The master keys should now be in sync. Check the logs to make sure no critical system logs related to the master key are present.
- 6. Enable configuration synchronize again by selecting the **Enable Config Sync** check box and then commit the configuration.

HA master key update (keys out of sync or expired):

- 1. If the master keys are out of sync or expired, you will see critical errors in the system log. If this occurs, you should immediately disable HA config sync on both devices in the HA pair. Navigate to **Device > High Availability** on both devices and from the **General** tab deselect the **Enable Config Sync** check box. Commit the configuration on both devices.
- 2. Ensure that no pending configuration updates exist on both devices in the HA pair. If there are pending changes, commit the configuration on both devices and wait until all configuration updates are complete. You can view the commit status by clicking the Tasks link on the bottom right of the web interface on each device. From the CLI, run show jobs all to see all running jobs or show jobs pending to view pending jobs. To check that no pending updates are present, from configure mode, run check pending-changes and you should see No.
- 3. Update the master key on device A using 16 characters and commit the configuration.
- 4. Update the master key on device B with the same master key and commit the configuration.
- 5. Enable configuration synchronization on both devices and commit the configuration.
- 6. The master keys should now be in sync. Check the logs to make sure no critical system logs related to the master key are present.
- 7. Enable configuration synchronize again by selecting the **Enable Config Sync** check box and then commit the configuration. If problems persist, please contact the Technical Support department.

High Availability

PAN-OS supports active/passive and active/active high availability (HA).



Note: In an HA pair, both firewalls must be the same model and have the same licenses. If state synchronization is enabled, existing sessions continue after a switchover; however, threat prevention functions do not continue. Threat protection will apply to new sessions.

Active/Passive HA

In the active/passive configuration, two devices form an HA group to provide redundancy. The two firewalls mirror each other in configuration. If the active firewall fails for any reason, the passive firewall becomes active automatically with no loss of service. A failover can also occur if selected Ethernet links fail or if the active firewall cannot reach one or more of the specified destinations. From a traffic processing perspective, at most one device receives packets at any one time.

The following rules apply to HA operation and failover:

- The active firewall continuously synchronizes its configuration and session information with the passive firewall over the HA interfaces.
- If the active firewall fails, then the passive firewall detects the loss of heartbeats and automatically becomes active.
- If one HA interface fails, synchronization continues over the remaining interface. If the state synchronization connection is lost, then no state synchronization occurs. If the configuration synchronization is lost, heartbeats are lost. Both devices determine that the other is down, and both become active.
- You can configure the management ports (MGT) on the HA devices to provide a backup path for heartbeat and hello messages. If you configure HA1 or HA1 backup on the management port, you do not have to enable the heartbeat backup option.

Active/Active HA

Active/active high availability allows both devices in an HA pair to pass traffic concurrently and is deployed primarily in asymmetrically routed environments where App-ID and Content-ID support are required. Layer 7 inspection for App-ID and Content-ID is performed on a single device for each session (that device is known as the session owner). PAN-OS uses packet forwarding (through the HA3 link), where required, to send packets to the designated session owner for processing.



Note: You must enable jumbo frames on the firewall and on all intermediary networking devices when using the HA3 interface. To enable jumbo frames, select **Device > Setup < Session** and select the option to **Enable Jumbo Frame** in the Session Settings section.

Active/active devices can be deployed with Layer 3 or virtual wire interfaces. In Layer 3 deployments, the scanned packets can be forwarded directly by the session owner after processing. In virtual wire deployments, the scanned packets must be returned to the receiving firewall to preserve the forwarding path. If the session owner receives the packet

initially, the HA3 link is not used. Sessions that do not require App-ID and Content-ID are forwarded directly by the receiving device (even if it is not the session owner) to maximize performance.

To provide flexibility, you can configure Layer 3 interfaces in several ways. Note that it often makes sense to configure Layer 3 interfaces with a static interface IP address in addition to a virtual address (floating IP address or ARP load sharing IP address).

- **Static interface IP**—Layer 3 interfaces should be assigned static IP addresses whenever the firewall will be participating in dynamic routing protocols with neighboring devices. One possible active/active deployment option makes use of dynamic routing protocol cost metrics to force a symmetric path through the HA pair. In this case, all traffic will be symmetric and the efficiency of the active/active pair will be maximized.
- **Floating IP**—This mode is employed when Virtual Router Redundancy Protocol (VRRP)like functionality is required, such as when an IP address must be available regardless of the state of the HA pair members. It is typical to configure two floating IP addresses on a particular interface such that each firewall owns one. Ownership is assigned to the device ID that has the higher priority. If either firewall fails, the floating IP address will be transitioned to the HA peer.
- **ARP load sharing**—This mode is used to distribute the load of host traffic between the two firewalls using Address Resolution Protocol (ARP).

For a more in-depth discussion of these three options refer to the discussions later in this section.

Packet Flow

Packet flow works as follows in an active/active configuration:

• The session owner is responsible for all packet processing for App-ID and Content-ID. The session owner can be configured to be (1) the first device that receives a packet for the session or (2) the primary device. If the configuration option is set to "primary device," all sessions are set up on the primary device.



Note: Logs passing through an active/active HA pair appear on the device that is designated as the session owner.

- A single device is selected as the session setup device for all new sessions. This is necessary to avoid possible race conditions that can occur in asymmetrically routed environments. The session setup device is determined by one of the following methods:
 - IP modulo—Uses a simple modulus operation on the source IP address to determine which device will set up the session. IP modulo distributes session setup responsibilities to a particular HA device according to the parity of the IP address.
 - Primary Device—Session setup always occurs on the primary device.
 - **Hash**—Hashing is used to inject more randomness in the setup device selection process.

- When a new session begins, the receiving firewall either sets up the session or forwards it to the HA peer. The action is determined by the session setup configuration, as described previously. During this time, the session owner (the device responsible for maintaining state for App-ID and Content-ID) is determined according to the configuration.
- If packets arrive at the session owner, the packet is scanned for threats (if configured in security policy) and forwarded according to the device's networking configuration. If packets arrive at the HA peer, a session table lookup identifies that the session is owned by the other device and the packet can be forwarded across HA3 to the session owner. If Layer 7 inspection is not required for the session, the receiving device can simply match the session with an existing session table entry and forward the packet towards its final destination.

Deployment Options

Active/active HA supports concurrent use of virtual wire and Layer 3 interfaces. All active/ active deployment options are supported in IPv6 environments, including IPv6 path monitoring.

Virtual Wire Deployment

Virtual wire deployments support full asymmetric routing as with other active/active deployments. It is important to note that packets forwarded to the session owner for App-ID and Content-ID inspection must be returned to the receiving firewall to preserve the forwarding path.

Layer 3-Floating IP Deployment

This deployment option allows for the creation of floating IP addresses that can move between the HA devices when a link failure or device failure occurs. The port that owns the floating IP address responds to ARP requests with a virtual MAC address. Floating IP addresses are recommended when VRRP-like functionality is required. Floating IP addresses can be used in VPN and Network Address Translation (NAT) configurations, allowing for persistent connections when a failure occurs on the device offering those services.



Note: Floating IP address will utilize a different virtual MAC address when the IP moves between HA devices when a failure occurs.

Layer 3-ARP Load-Sharing

ARP load-sharing allows the HA pair to share an IP address and provide gateway services. In this scenario, all hosts are configured with a single gateway IP address. ARP requests for the gateway IP address are responded to by a single device in the pair, according to the source of the ARP request. The device selection algorithm can be tuned to achieve a more even distribution of host traffic between the two firewalls. ARP load-sharing should be used when the firewall and hosts exist on the same broadcast domain. If Layer 3 separation exists, the benefits of ARP load-sharing will be lost.



Note: You cannot ping or perform any management services on an ARP loadsharing IP address in active/active mode.

Layer 3-Route Based Redundancy (Static Interface IPs)

Route based redundancy forces traffic to be symmetric by using routing metrics such as Open Shortest Path First (OSPF) costs on the firewalls and on neighboring devices. Load sharing can be handled by adjusting costs to route traffic through both firewalls. In this case, the IP address assigned to the device interface is pinned down and does not fail over to the HA peer during a failover.

NAT Considerations

In active/active mode, it is necessary to define an active/active device binding in all NAT rules. Active/active device binding becomes available in the web interface when the HA mode has been changed to active/active. When a new session is created, device binding determines which NAT rules are matched by the firewall (the device binding must include the session owner device to produce a match). Although NAT policy match is performed by the session setup device, NAT rules are evaluated from the perspective of the session owner. The session is translated according to NAT rules that are bound to the session owner device. For device-specific rules, a firewall skips all NAT rules that are not bound to the session owner when the NAT policy match is performed.

For example, suppose device 1 is the session owner and is also responsible for setting up the session. When device 1 attempts to match the session to a NAT rule, it will skip all rules with a device binding of device 0.

NAT device binding options include the following:

- **Device 0 and Device 1**—Translation is performed according to device-specific bindings only if the session owner and the device ID in the NAT rule match. Device-specific NAT rules are commonly used when the two firewalls use unique public IP addresses for translation.
- **Both**—This option allows either device to match new sessions to the NAT rule and is commonly used for destination NAT.
- **Primary**—This option allows only the active-primary device to match new sessions to the NAT rule. This setting is used mainly for inbound static NAT, where only one firewall should respond to ARP requests. Unlike device 0/1 bindings, a primary device binding can move between devices when the primary role is transferred.

The following scenarios apply to active/active NAT deployments.

Source Translation to Dynamic IP or IP/Port Pool

When source translating to a dynamic IP or dynamic IP/port pool, it is necessary to tie NAT rules to a specific device (Device ID 0 or 1). The IP pools to which the HA devices are translating must not overlap. When a session is established, either device can translate return packets.

In the following example, the sessions owned by device 0 are translated to 1.1.1.1 and the sessions owned by device 1 are translated to 1.1.1.2. In the event of a device failure, sessions from device 0 will continue to be translated to 1.1.1.1 until they have ceased. In this example, it is valid to use floating IP addresses on each of the firewalls if that functionality is required.



Figure 3. Source Translation Dynamic IP Configuration

	Original P	Original Packet		
Name	e Source Destination Zone Zone		Source Active/Active H Translation Binding	
Src NAT Device 0	L3Trust	L3Untrust	dynamic-ip-and- port 1.1.1.1	0
Src NAT Device 1	L3Trust	L3Untrust	dynamic-ip-and- port 1.1.1.2	1

Table 40. Source Translation Dynamic IP Rules

Dynamic Source Translation to Public IP Addresses for Different Internet Service Providers (ISPs)

In this scenario, NAT rules are tied to specific devices (Device ID 0 or 1). All sessions owned by device 0 are translated to 1.1.1.1 and all sessions owned by device 1 are translated to 2.2.2.1. If device 0 fails, device 1 will attempt to translate existing sessions according to the original IP address of 1.1.1.1. If the second ISP cannot route to these addresses, the sessions will fail. In this example, the ISP-specific interface IP addresses are pinned down to a particular device. A floating IP address should not be used in this configuration.



Figure 4. Dynamic Source Translation to Public IP Address Configuration

	Original Packet		Translated Packet	
Name	lame Source Destination Zone Zone		Source Translation	Active/Active HA Binding
Src NAT Device 0	L3Trust	L3Untrust	dynamic-ip-and- port 1.1.1.1	0
Src NAT Device 1	L3Trust	L3Untrust	dynamic-ip-and- port 2.2.2.1	1

Table 41. Dynamic Source Translation to Public IP Address Rules

Destination Translation to Provider-Independent IP Address

In this scenario, NAT rules are tied to both devices. Translation is the same regardless of which device receives the first incoming packet. A packet destined to 3.3.3.30 will be translated to 10.0.0.200 regardless of which device receives the packet.



Figure 5. Destination Translation to Provider-Independent IP Address

Table 42.	Dynamic	Source	Translation t	o Public IP	Address R	ules
-----------	---------	--------	----------------------	-------------	-----------	------

		Original Packet		Translated Packet	
Name	Source Zone	ne Destination Destination Zone Address		Destination Translation	Active/Active HA Binding
DNAT Prov Indep	L3Untrust	L3Untrust	3.3.3.30	address: 10.0.0.200	both

Setting Up HA

To set up HA, follow these steps:

- 1. Use two firewalls with the same model number.
- 2. Mount both firewalls in a rack near each other, and power them up as described in the *Hardware Reference Guide*. If you are enabling HA on an existing production device, backup the configuration and you may also want to perform a factory reset on the new device that will be introduced. This will ensure that the configuration is clean on the new device. After HA is configured, you would then push the configuration from the existing primary device to the new secondary device. Make sure that you do not push the factory config from the new device to the existing device. For information on performing a factory reset, refer to the *PAN-OS Command Line Interface Reference Guide*.
- 3. Connect each firewall to your network and the Internet using the same physical ports.

4. Using two crossover RJ-45 Ethernet cables, connect the HA1 and HA2 ports on each firewall to the same ports on the other firewall, or connect the ports on both firewalls to a switch. HA1 is for the control link, and HA2 is for the data link. For active/active configurations, make an additional physical connection, HA3, between the two firewalls. Link aggregation groups are recommended for link redundancy on HA3 when the firewall supports aggregate Ethernet.



Note: For devices that do not have dedicated HA interfaces, you must use the traffic interfaces for HA. For example, connect the ethernet 1/15 interfaces to each other and the ethernet1/16 interfaces to each other.

5. Open the **Network** tab and verify that the HA links are up. Configure each to be of the type HA.

interfaces							
	Interface	Interface Type	Management Profile	Link State	IP Address		
Δ	ethernet1/15	HA					
A	ethernet1/16	НА					

Figure 6. Verifying HA Interfaces

6. Configure HA settings on both firewalls. Refer to "Enabling HA on the Firewall" on page 101.

Item to note when setting up HA

Crossover cables are recommended when HA links are directly connected.

Enabling HA on the Firewall

Device > High Availability

After setting up HA as described in "Setting Up HA" on page 99, you can enable HA on both the active and passive firewall. For each section on the **High Availability** page, click **Edit** in the header, and specify the corresponding information described in the following table.

Field	Description
General Tab	
Setup	Specify the following settings:
	• Enable HA—Activate HA functionality.
	• Group ID —Enter a number to identify the active/passive pair (1 to 63). Allows multiple pairs of active/passive firewalls to reside on the same network. The ID must be unique when more than one high availability pair resides on an Layer 2 network.
	• Description —Enter a description of the active/passive pair (optional).
	• Mode—Choose active-active or active-passive.
	• Peer HA IP Address —Enter the IP address of the HA1 interface that is specified in the Control Link section of the other firewall.
	• Backup Peer HA IP Address—Enter the IP address for the peer's backup control link.
	• Enable Config Sync—Synchronize the peer system.
	• Link Speed— Select the speed for the data link between the active and passive firewalls (Firewalls with dedicated HA ports).
	• Link Duplex— Select a duplex option for the data link between the active and passive firewalls (Firewalls with dedicated HA ports).

Table 43. HA Settings

Field	Description
Election Settings	Specify the following settings:
	• Device Priority —Enter a priority value to identify the active firewall. The firewall with the lower value (higher priority) becomes the active firewall (range 0-255).
	• Heartbeat Backup—Uses the management ports on the HA devices to provide a backup path for heartbeat and hello messages. The managemer port IP address will be shared with the HA peer through the HA1 control link. No additional configuration is required.
	• Preemptive —Enable the higher priority firewall to resume active operation after recovering from a failure. If this setting is off, then the lower priority firewall remains active even after the higher priority firewall recovers from a failure.
	• Preemption Hold Time —Enter the time a passive or active-secondary device will wait before taking over as the active or active-primary device (range 1-60 min, default 1 min).
	• Promotion Hold Time —Enter the time that the passive device (in active, passive mode) or the active-secondary device (in active/active mode) will wait before taking over as the active or active-primary device after comm nications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made.
	• Hello Interval —Enter the number of milliseconds between the hello packets sent to verify that the HA program on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms.
	• Heartbeat Interval—Specify how frequently the HA peers exchange hear beat messages in the form of an ICMP ping (range 1000-60000 ms, default 1000 ms).
	• Maximum No. of Flaps—A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. You can specify the maximum number of flaps that are permitted before the firewar is determined to be suspended and the passive firewall takes over (range 16, default 3). The value 0 means there is no maximum (an infinite number of flaps is required before the passive firewall takes over).
	• Monitor Fail Hold Up Time (ms)—Specify the interval during which the firewall will remain active following a path monitor or link monitor failur. This setting is recommended to avoid an HA failover due to the occasiona flapping of neighboring devices (range 0-60000 ms, default 0 ms).
	• Additional Master Hold Up Time (min)—This time interval is applied to the same event as Monitor Fail Hold Up Time (range 0-60000 ms, default 500 ms). The additional time interval is applied only to the active device i active/passive mode and to the active-primary device in active/active mode. This timer is recommended to avoid a failover when both devices experience the same link/path monitor failure simultaneously

Table 43. HA Settings (Continued)

Field	Description
Control Link (HA1)/ Control Link (HA1 Backup)	The recommended configuration for the HA control link connection is to use the dedicated HA1 link between the two devices and use the management port as the Control Link (HA Backup) interface. In this case, you do not need to enable the Heartbeat Backup option in the Elections Settings page. If you are using a physical HA1 port for the Control Link HA link and a data port for Control Link (HA Backup), it is recommended that enable the Heartbeat Backup option.
	For devices that do not have a dedicated HA port, such as the PA-200, you should configure the management port for the Control Link HA connection and a data port interface configured with type HA for the Control Link HA1 Backup connection. Since the management port is being used in this case, there is no need to enable the Heartbeat Backup option in the Elections Settings page because the heartbeat backups will already occur through the management interface connection.
	Note: When using a data port for the HA control link, you should be aware that since the control messages have to communicate from the dataplane to the management plane, if a failure occurs in the dataplane, HA control link information cannot communicate between devices and a failover will occur. It is best to use the dedicated HA ports, or on devices that do not have a dedicated HA port, use the management port.
	 Specify the following settings for the primary and backup HA control links: Port—Select the HA port for the primary and backup HA1 interfaces. The backup setting is optional.
	Note: The management port can also be used as the control link.
	• IPv4/IPv6 Address —Enter the IPv4 or IPv6 address of the HA1 interface for the primary and backup HA1 interfaces. The backup setting is optional.
	• Netmask—Enter the network mask for the IP address (such as "255.255.255.0") for the primary and backup HA1 interfaces. The backup setting is optional.
	• Gateway —Enter the IP address of the default gateway for the primary and backup HA1 interfaces. The backup setting is optional.
	• Link Speed (Models with dedicated HA ports only)—Select the speed for the control link between the firewalls for the dedicated HA1 port.
	• Link Duplex (Models with dedicated HA ports only)—Select a duplex option for the control link between the firewalls for the dedicated HA1 port.
	• Encryption Enabled—Enable encryption after exporting the HA key from the HA peer and importing it onto this device. The HA key on this device must also be exported from this device and imported on the HA peer. Con- figure this setting for the primary HA1 interface. The key import/export is done on the Certificates page. Refer to "Importing, Exporting and Generating Security Certificates" on page 60.
	• Monitor Hold Time (ms)—Enter the length of time (milliseconds) that the firewall will wait before declaring a peer failure due to a control link failure (1000-60000 ms, default 3000 ms). This option monitors the physical link status of the HA1 port(s).

Table 43. HA Settings (Continued)

Field	Description
Data Link (HA2)	Specify the following settings for the primary and backup data link:
	 Port—Select the HA port. Configure this setting for the primary and backup HA2 interfaces. The backup setting is optional.
	• IP Address —Specify the IPv4 or IPv6 address of the HA interface for the primary and backup HA2 interfaces. The backup setting is optional.
	 Netmask—Specify the network mask for the HA interface for the primary and backup HA2 interfaces. The backup setting is optional.
	• Gateway —Specify the default gateway for the HA interface for the primary and backup HA2 interfaces. The backup setting is optional. If the HA2 IP addresses of the firewalls in the HA pair are in the same subnet, the Gateway field should be left blank.
	• Enable Session Synchronization—Enable synchronization of the session information with the passive firewall, and choose a transport option.
	• Transport—Choose one of the following transport options:
	 Ethernet—Use when the firewalls are connected back-to-back or through a switch (Ethertype 0x7261).
	- IP —Use when Layer 3 transport is required (IP protocol number 99).
	 UDP—Use to take advantage of the fact that the checksum is calculate on the entire packet rather than just the header, as in the IP option (UD port 29281).
	• Link Speed (Models with dedicated HA ports only)—Select the speed for the control link between the active and passive firewalls for the dedicated HA2 port.
	• Link Duplex (Models with dedicated HA ports only)—Select a duplex option for the control link between the active and passive firewalls for th dedicated HA2 port.
	• HA2 keep-alive—Select this check box to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set, the defined action will occur (log or split data-path). The optic is disabled by default. You can configure the HA2 keep-alive option on both devices, or just one device in the HA pair. If the option is only set on one device, only that davias will can d the keep alive measures. The other device will be patient
	though if a failure occurs and will go into split data path mode if that actions selected.
	 Action—Select the action to take if monitoring messages fail based on the threshold setting.
	 log-only—Select to generate a critical level system log message when an HA2 failure occurs based on the threshold setting. If the HA2 path recovers, an informational log will be generated. In an active/passive configuration, you should use this action since there's no need to split the data since only one device is active at a given time.
	> split data-path—This action is designed for an active/active H configuration. In an <i>active/active</i> configuration, if session synchronization is disabled by the admin, or by a monitoring failure, session ownership and session setup will both be set to the local device and new sessions will be processed locally for th session lifetime

Table 43. HA Settings (Continued)

Field	Description
	 Threshold (ms)—The duration in which keep-alive messages have failed before one of the above actions will be triggered (range 5000-60000ms, default 10000ms).
	Note: When an HA2 backup link is configured, failover to the backup link will occur if there is a physical link failure. With the HA2 keep-alive option enabled, the failover will also occur if the HA keep-alive messages fail based on the defined threshold.
Link and Path Mo	nitoring Tab
Path Monitoring	Specify the following:
	• Enabled —Enable path monitoring. Path monitoring enables the firewall to monitor specified destination IP addresses by sending ICMP ping messages to make sure that they are responsive. Use path monitoring for virtual wire, Layer 2, or Layer 3 configurations where monitoring of other network devices is required for failover and link monitoring alone is not sufficient.
	• Failure Condition—Select whether a failover occurs when any or all of the monitored path groups fail to respond.
Path Group	Define one or more path groups to monitor specific destination addresses. To add a path group, click Add for the interface type (Virtual Wire, VLAN, or Virtual Router) and specify the following:
	• Name —Enter a name to identify the group.
	• Enabled —Enable the path group.
	• Failure Condition—Select whether a failure occurs when any or all of the specified destination addresses fails to respond.
	• Source IP —For virtual wire and VLAN interfaces, enter the source IP address used in the probe packets sent to the next-hop router (Destination IP address). The local router must be able to route the address to the firewall. The source IP address for path groups associated with virtual routers will be automatically configured as the interface IP address that is indicated in the route table as the egress interface for the specified destination IP address.
	• Destination IPs —Enter one or more (comma-separated) destination addresses to be monitored.
	• Ping Interval —Specify the interval between pings that are sent to the destination address (range 200-60,000 milliseconds, default 200 milliseconds).
	• Ping Count —Specify the number of failed pings before declaring a failure (range 3-10 pings, default 10 pings).

Table 43. HA Settings (Continued)

Field	Description
Link Monitoring	Specify the following:
	• Enabled —Enable link monitoring. Link monitoring allows failover to be triggered when a physical link or group of physical links fails.
	• Failure Condition—Select whether a failover occurs when any or all of the monitored link groups fail.
Link Groups	Define one or more link groups to monitor specific Ethernet links. To add a link group, specify the following and click Add :
	• Name—Enter a link group name.
	• Enabled —Enable the link group.
	• Failure Condition—Select whether a failure occurs when any or all of the selected links fail.
	• Interfaces—Select one or more Ethernet interfaces to be monitored.
Active/Passive Tab	
Passive Link State	Choose from the following options:
	• auto —Causes the link status to reflect physical connectivity, but discards all packets received. This option allows the link state of the interface to stay up until a failover occurs, decreasing the amount of time it takes for the passive device to take over. This option is supported in Layer 2, Layer 3, and Virtual Wire mode. The auto option is desirable, if it is feasible for your network.
	Note: When set to auto, the link state on the passive peer in a VM-Series firewall configured as an HA pair displays in a down state on Network > Interfaces > Ethernet . The link state for the interface icon is red.
	• shutdown —Forces the interface link to the down state. This is the default option, which ensures that loops are not created in the network.
Monitor Fail Hold Down Time	Specify the length of time (minutes) that a firewall will spend in the non- functional state before becoming passive. This timer is used only when the failure reason is a link or path monitor failure (range 1 to 60, default 1).
Active/Active Confi	ig Tab
Packet Forwarding	Select the Enable check box to enable forwarding of packets over the HA3 link. This is required for asymmetrically routed sessions that require Layer 7 inspection for App-ID and Content-ID.
HA3 Interface	Choose the interface to forward packets between HA peers when configured in active/active mode.
	Note: You must enable jumbo frames on the firewall and on all intermediary networking devices when using the HA3 interface. To enable jumbo frames, select Device > Setup < Session and select the option to Enable Jumbo Frame in the Session Settings section.
VR Sync	Force synchronization of all virtual routers configured on the HA devices.
	Virtual Router synchronization can be used when the virtual router is not employing dynamic routing protocols. Both devices must be connected to the same next-hop router through a switched network and must use only static routing.

Table 43. HA Settings (Continued)

Field	Description
QoS Sync	Synchronize the QoS profile selection on all physical interfaces. Use this option when both devices have similar link speeds and require the same QoS profiles on all physical interfaces. This setting affects the synchronization of QoS settings on the Network tab. QoS policy is synchronized regardless of this setting.
Tentative Hold Time (sec)	When a firewall in an HA active/active state fails it will go into a tentative state. This timer defines how long it will stay in this state. During the tentative period the firewall will attempt to build routing adjacencies and populate its route table before it will process any packets. Without this timer, the recovering firewall would enter the active-secondary state immediately and would blackhole packets because it would not have the necessary routes (default 60 seconds).
Session Owner	Specify one of the following options for selecting the session owner:
Selection	• Primary Device —Select this option to have the active/primary firewall handle Layer 7 inspection for all sessions. This setting is recommended primarily for troubleshooting operations.
	• First packet—Select this option to make the firewall that receives the first packet of the session responsible for Layer 7 inspection in support of App-ID and Content-ID. This is the recommended configuration to minimize utilization of the HA3 packet forwarding link.
Session Setup	Choose the method for initial session setup.
-	• IP Modulo—Selects a firewall based on the parity of the source IP address.
	• Primary Device —Ensures that all sessions are set up on the primary fire- wall.
	• IP Hash —Determines the setup firewall using a hash of the source IP address or source and destination IP address, and hash seed value if more randomization is desired.

Table 43. HA Settings (Continued)

Field	Description
Virtual Address	Click Add , select the IPv4 or IPv6 tab and then click Add again to enter options for an HA virtual Address that will be used by the HA active/active cluster. You can select the type of virtual address to be either Floating or ARE Load Sharing. You can also mix the type of virtual address types in the cluster, for example, you could use ARP load sharing on the LAN interface and a Floating IP on the WAN interface. For more information deployment types, refer to "Deployment Options" on page 95.
	• Floating —Enter an IP address that will move between HA devices in the event of a link or device failure. You should configure two floating IP addresses on the interface, so that each firewall will own one and then set the priority. If either firewall fails, the floating IP address will be transitioned to the HA peer.
	 Device 0 Priority—Set the priority to determine which device will own the floating IP address. A device with the lowest value will have the highest priority.
	 Device 1 Priority—Set the priority to determine which device will own the floating IP address. A device with the lowest value will have the highest priority.
	 Failover address if link state is down—Use the failover address when the link state is down on the interface.
	• ARP Load Sharing —Enter an IP address that will be shared by the HA pair and will provide gateway services for hosts. This option should only be used when the firewall and hosts exist on the same broadcast domain. Select the Device Selection Algorithm :
	 IP Modulo—If this option is selected, the firewall that will respond to ARP requests will be selected based on the parity of the ARP requesters IP address.
	 IP Hash—If this option is selected, the firewall that will respond to ARP requests will be selected based on a hash of the ARP requesters IP address.
Operational Comm	ands
Suspend local device	Put the HA device in suspend mode, which temporarily disables HA functionality on the firewall. If you suspend the current active firewall, the

Table 43. HA Settings (Continued)

Important items to consider when configuring HA

• The firewall that is assigned the lower device priority value is the higher priority device and becomes the active firewall in an HA pair when preemption is enabled on both firewalls in the pair.

secondary firewall will take over.

- The Preemption option must be enabled on both devices for the higher priority firewall to resume active operation upon recovery following a failure.
- The subnet that is used for the local and peer IP should not be used anywhere else on the virtual router.
- The OS and Content versions should be the same on each device. A mismatch can prevent the devices in the cluster from synchronizing.
- The LEDs are green on the HA ports for the active firewall and amber on the passive firewall.
- To test failover, pull a cable on the active device, or put the active device into a suspend state by issuing the CLI command request high-availability state suspend. You can also suspend the active device by clicking the **Suspend** link at **Device > High** Availability > Operational Commands tab.
- To place a suspended device back into a functional state, use the CLI command request high-availability state functional.
- To view detailed HA information about the local firewall, use the CLI command **show** high-availability all.
- To compare the configuration of the local and peer firewalls, use the CLI command **show high-availability state** from either device. You can also compare the configurations on the local and peer firewalls using the **Config Audit** tool on the **Device** tab by selecting the desired local configuration in the left selection box and the peer configuration in the right selection box.
- Synchronize the firewalls from the web interface by pressing the **Push Configuration** button located in the HA widget on the **Dashboard** tab. Note that the configuration on the device from which you push the configuration overwrites the configuration on the peer device. To synchronize the firewalls from the CLI on the active device, use the command request high-availability sync-to-remote running-config.
- To follow the status of the load, use the CLI command **show jobs processed**.



Note: In a High Availability (HA) active/passive configuration with devices that use 10 Gigabit SFP+ ports, when a failover occurs and the active device changes to a passive state, the 10 Gigabit Ethernet port is taken down and then brought back up to refresh the port, but does not enable transmit until the device becomes active again. If you have monitoring software on the neighboring device, it will see the port as flapping because it is going down and then up again. This is different behavior than the action with other ports, such as the 1 Gigabit Ethernet port, which is disabled and still allows transmit, so flapping is not detected by the neighboring device.

HA Lite

The PA-200 and VM-Series firewalls supports a "lite" version of active/passive HA that does not include any session synchronization. HA lite does provide configuration synchronization and synchronization of some runtime items. It also supports failover of IPSec tunnels (sessions must be re-established), DHCP server lease information, DHCP client lease information, PPPoE lease information, and the firewall's forwarding table when configured in Layer 3 mode.

Virtual Systems

A virtual system specifies a collection of physical and logical firewall interfaces (including VLANs, and virtual wires) and security zones. (For more information on security zones, refer to "Defining Security Zones" on page 151.) Virtual systems allow you to segment the administration of all policies (security, NAT, QoS, etc.) as well as all reporting and visibility functions provided by the firewall.

Virtual systems generally operate on the security functionality of the firewall. Networking functions including static and dynamic routing are not controlled by virtual systems. If routing segmentation is desired for each virtual system, you must create an additional virtual router.



Note: The PA-4000 and PA-5000 Series firewalls support multiple virtual systems. The PA-2000 and PA-3000 Series firewalls can support multiple virtual systems if the appropriate license is installed. The PA-500 and PA-200 firewalls do not support virtual systems.

For example, if you want to customize the security features for the traffic that is associated with your Finance department, you can define a Finance virtual system and then define security policies to apply only to that department.

Figure 7 illustrates the relationship between policies and virtual systems in the firewall. Policies are associated with individual virtual systems, by contrast with device and network level functions, which apply to the overall firewall.



Figure 7. Virtual Systems and Policies

To optimize policy administration, you can create virtual system administrator accounts that allow access to individual virtual systems, while maintaining separate administrator accounts for overall device and network functions. For example, a virtual system administrator in the Finance department can be assigned to manage the security policies only for that department. Initially all interfaces, zones, and policies belong to the default virtual system (vsys1). When you enable multiple virtual systems, note the following:

- All items needed for policies are created and administered by a virtual systems administrator.
- Zones are objects within virtual systems. Before defining a policy or policy object, select the virtual system from the **Virtual System** drop-down list on the **Policies** or **Objects** tab.
- Interfaces, VLANs, virtual wires, and virtual routers can be assigned to virtual systems. Refer to "Defining Virtual Systems" on page 113.
- Remote logging destinations (SNMP, syslog, and email), as well as applications, services, and profiles, can be shared by all virtual systems or be limited to a selected virtual system.

Communications Among Virtual Systems

The virtual systems in the firewall are treated as separate entities. To support internal traffic flows between virtual systems, you must indicate which virtual systems are able to communicate with each other. You do so when configuring a virtual system by specifying the other virtual systems that are visible to it. When the virtual systems are made visible to each other, create "external"-type zones and specify which virtual systems will map to each external zone. In the following example, Dept 1 VSYS must have an external zone that refers to Dept 2 VSYS for use in all policies affecting traffic passing between the virtual systems. Traffic log entries are recorded in both virtual systems for inter-VSYS traffic.

Each virtual system must have policies for sending and receiving traffic. For example, allowing Dept 1 VSYS to communicate with Dept 2 VSYS requires a policy in Dept 1 VSYS to allow traffic to go to Dept 2 VSYS and a policy in Dept 2 VSYS to accept incoming traffic from Dept 1 VSYS.



Figure 8. Communications Among Virtual Systems

Shared Gateways

In a standard virtual system interface configuration, each virtual system uses a dedicated interface to the outside world. Each virtual system is autonomous, and there are no direct communication paths among the virtual systems that are internal to the firewall, unless such communications are explicitly configured (refer to "Communications Among Virtual Systems" on page 111). Because each virtual system has its own IP address, multiple addresses are required for external communications.



Figure 9. Virtual Systems Without a Shared Gateway

Shared gateways allow virtual systems to share a common interface for external communications. This is especially helpful in deployments where the ISP provides only a single IP address. All of the virtual systems communicate with the outside world through the physical interface using a single IP address (see Figure 10). A single virtual router is used to route the traffic for all of the virtual systems through the shared gateway.



Figure 10. Virtual Systems with a Shared Gateway

All policy rules are managed at the virtual system level. You can create NAT and policy-based forwarding rules through the shared gateway, if needed, by selecting the shared gateway from the Virtual System drop-down list on the policy screen.

Defining Virtual Systems

Device > Virtual Systems

To define virtual systems, you must first enable the definition of multiple virtual systems. To do so, open the **Device > Setup** page, click the **Edit** link under **General Settings** in the **Management** tab, and select the **Multi Virtual System Capability** check box. This adds a **Virtual Systems** link to the side menu.

You can now open the **Virtual Systems** page, click **Add**, and specify the following information.

Field	Description
ID	Enter an integer identifier for the virtual system. Refer to the data sheet for your firewall model for information on the number of supported virtual systems.
Name	Enter a name (up to 31 characters) to identify the virtual system. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required.
General Tab	Select a DNS proxy profile from the drop-down list if you want to apply DNS proxy rules to this interface. Refer to "DNS Proxy" on page 173.
	To include objects of a particular type, select the check box for that type (interface, VLAN, virtual wire, virtual router, or visible virtual system). Click Add and choose from the drop-down list. You can add one or more objects of any type. To remove an object, select it and click Delete .

Table 44. Virtual System Settings

Field	Description
Resource Tab	Enter the following settings:
	• Sessions Limit—Maximum number of sessions allowed for this virtual system.
	• Security Rules—Maximum number of security rules allowed for this virtual system.
	• NAT Rules—Maximum number of NAT rules allowed for this virtual system.
	 Decryption Rules—Maximum number decryption rules allowed for this virtual system.
	• QoS Rules —Maximum number of QoS rules allowed for this virtual system.
	• Application Override Rules—Maximum number of application over- ride rules allowed for this virtual system.
	• PBF Rules —Maximum number of policy based forwarding (PBF) rules allowed for this virtual system.
	• CP Rules —Maximum number of captive portal (CP) rules allowed for this virtual system.
	• DoS Rules — Maximum number of denial of service (DoS) rules allowed for this virtual system.
	• Site to Site VPN Tunnels—Maximum number of site-to-site VPN tunnels allowed for this virtual system.
	• Concurrent GlobalProtect Tunnel Mode Users—Maximum number of concurrent remote GlobalProtect users allowed for this virtual system.

Table 44. Virtual System Settings (Continued)

After defining the virtual systems, you can perform any of the following additional tasks:

- To change a virtual system, click the virtual system name or the name of the interface, VLAN, virtual wire, virtual router, or visible virtual systems you want to change, make the appropriate changes, and click **OK**.
- To define security zones for the new virtual system, choose **Network > Zones** and define security zones for each new virtual system (refer to "Defining Security Zones" on page 151). When you define a new zone, you can now select a virtual system.
- Click **Network > Interfaces** and verify that each interface has a virtual system and security zone.

Configuring Shared Gateways

Device > Shared Gateways

Shared gateways use Layer 3 interfaces, and at least one Layer 3 interface must be configured as a shared gateway. Refer to "Configuring Layer 3 Interfaces" on page 130.

Table 45. Shared Gateway Settings

Field	Description
ID	Identifier for the gateway (not used by firewall).
Name	Enter a name for the shared gateway (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required.
Interfaces	Select check boxes for the interfaces that the shared gateway will use.

Defining Custom Response Pages

Device > Response Pages

Custom response pages are the web pages that are displayed when a user tries to access a URL. You can provide a custom HTML message that is downloaded and displayed instead of the requested web page or file.

Each virtual system can have its own custom response pages.

The following table describes the types of custom response pages that support customer messages.



Note: Refer to Appendix A, "Custom Pages" for examples of the default response pages.

Page Type	Description
Antivirus Block	Access blocked due to a virus infection.
Application Block	Access blocked because the application is blocked by a security policy.
Captive Portal Comfort	Page for users to verify their user name and password for machines that are not part of the domain.
File Blocking Block	Access blocked because access to the file is blocked.
File Blocking Continue	Page for users to confirm that downloading should continue. This option is available only if continue functionality is enabled in the security profile. Refer to "File Blocking Profiles" on page 220.
GlobalProtect Portal Help	Custom help page for GlobalProtect users (accessible from the portal).
GlobalProtect Portal Login	Page for users who attempt to access the GlobalProtect portal. For information on GlobalProtect, refer to "Overview" on page 335.

Table 46. Custom Response Page Types

Page Type	Description
GlobalProtect Welcome Page	Welcome page for users who attempt to log in to the GlobalProtect portal. For information on GlobalProtect, refer to "Overview" on page 335.
SSL Certificate Errors Notify Page	Notification that an SSL certificate has been revoked.
SSL Decryption Opt-out Page	User warning page indicating that this session will be inspected.
URL Filtering Continue and Override Page	 Page with initial block policy that allows users to bypass the block. For example, a user who thinks the page was blocked inappropriately can click the Continue button to proceed to the page. With the override page, a password is required for the user to override the policy that blocks this URL. See the "URL Admin Override" section of Table 1 for instructions on setting the override password.
URL Filtering and Category Match Block Page	Access blocked by a URL filtering profile or because the URL category is blocked by a security policy.

Table 46. Custom Response Page Types (Continued)

You can perform any of the following functions under **Response Pages**.

- To import a custom HTML response page, click the link of the page type you would like to change and then click import/export. Browse to locate the page. A message is displayed to indicate whether the import succeeded. For the import to be successful, the file must be in HTML format.
- To export a custom HTML response page, click the **Export** link for the type of page. Select whether to open the file or save it to disk, and select the check box if you want to always use the same option.
- To enable or disable the **Application Block** page or **SSL Decryption Opt-out** pages, click the **Enable** link for the type of page. Select or deselect the **Enable** check box.
- To use the default response page instead of a previously uploaded custom page, delete the custom block page and commit. This will set the default block page as the new active page.

Viewing Support Information

Device > Support

The support page allows you to access support related options. You can view the Palo Alto Networks contact information, view your support expiration date, and view product and security alerts from Palo Alto Networks based on the serial number of your firewall.

Perform any of the following functions on this page:

- **Support**—Use this section to view Palo Alto Networks support contact information, view support status for the device or activate your contract using an authorization code.
- **Production Alerts/Application and Threat Alerts**—These alerts will be retrieved from the Palo Alto Networks update servers when this page is accessed/refreshed. To view the details of production alerts, or application and threat alerts, click the alert name. Production alerts will be posted if there is a large scale recall or urgent issue related to a given release. The application and threat alerts will be posted if significant threats are discovered.
- **Links**—This section provides a link to the Support home page, the site to manage your cases, and a link to register the device using your support login.
- **Tech Support File**—Use the **Generate Tech Support File** link to generate a system file that the Support group can use to help troubleshoot issues that you may be experiencing with the device. After you generate the file, click **Download Tech Support File** to retrieve it and then send it to the Palo Alto Networks Support department.
- Stats Dump File—Use the Generate Stats Dump File link to generate a set of XML reports that summarizes network traffic over the last 7 days. Once the report is generated, click the Download Stats Dump File link to retrieve the report. The files are then used by a Palo Alto Networks or Authorized Partner systems engineer to generate an Application Visibility and Risk Report (AVR Report). The AVR highlights what has been found on the network and the associated business or security risks that may be present and is typically used as part of the evaluation process. For more information on the AVR Report, please contact you Palo Alto Networks or Authorized Partner systems engineer.

Viewing Support Information

Chapter 4 Network Configuration

This chapter describes how to configure the firewall to support your network architecture:

- "Firewall Deployment" in the next section
- "Firewall Interfaces" on page 127
- "Security Zones" on page 151
- "VLAN Support" on page 152
- "Virtual Routers and Routing Protocols" on page 153
- "DHCP Server and Relay" on page 171
- "DNS Proxy" on page 173
- "Network Profiles" on page 174



Note: For information about VPN support on the firewall, refer to "Configuring IPSec Tunnels" on page 309 and "Configuring IPSec Tunnels" on page 309. For information about quality of service (QoS) support, refer to "Configuring Quality of Service" on page 355

Firewall Deployment

The firewall can replace your existing firewall when installed between an edge router (or other device that faces the Internet) and a switch or router that connects to your internal network. The firewall supports a wide range of deployment options and interface types that can be used simultaneously on different physical interfaces. They are described in the following sections:

- "Virtual Wire Deployments" in the next section
- "Layer 2 Deployments" on page 124
- "Layer 3 Deployments" on page 124
- "Tap Mode Deployments" on page 125
- "Defining Virtual Wires" on page 125
- "Packet Content Modification" on page 126

Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together (Figure 11), and should be used only when no switching or routing is needed.

A virtual wire deployment allows the following conveniences:

- Simplifies installation and configuration.
- Does not require any configuration changes to surrounding or adjacent network devices.

The "default-vwire" that is shipped as the factory default configuration, binds together Ethernet ports 1 and 2 and allows all untagged traffic. You can, however, use a virtual wire to connect any two ports and configure it to block or allow traffic based on the virtual LAN (VLAN) tags; the VLAN tag "0" indicates untagged traffic. You can also create multiple subinterfaces, add them into different zones and then classify traffic according to a VLAN tag, or a combination of a VLAN tag with IP classifiers (address, range, or subnet) to apply granular policy control for specific VLAN tags or for VLAN tags from a specific source IP address, range, or subnet.



Figure 11. Virtual Wire Deployment

Virtual Wire Subinterfaces

Virtual wire subinterfaces provide flexibility in enforcing distinct policies when you need to manage traffic from multiple customer networks. It allows you to separate and classify traffic into different zones (the zones can belong to separate virtual systems, if required) using the following criteria:

- VLAN tags —The example in Figure 12, shows an Internet Service Provider (ISP) using virtual wire subinterfaces with VLAN tags to separate traffic for two different customers.
- VLAN tags in conjunction with IP classifiers (address, range, or subnet). The example in Figure 13, shows an Internet Service Provider (ISP) with two separate virtual systems on a firewall that manages traffic from two different customers. On each virtual system, the example illustrates how virtual wire subinterfaces with VLAN tags and IP classifiers are used to classify traffic into separate zones and apply relevant policy for customers from each network.

The general workflow for the configuration is:

- 1. Configure two Ethernet interfaces as type virtual wire, and assign these interfaces to a virtual wire. To set up virtual wires, see "Configuring Virtual Wire Interfaces" on page 138.
- Create subinterfaces on the parent Virtual Wire to separate CustomerA and CustomerB traffic. Make sure that the VLAN tags defined on each pair of subinterfaces that are configured as virtual wire(s) are identical. This is essential because a virtual wire does not switch VLAN tags.
 To set up virtual wire subinterfaces, see "Configuring Virtual Wire Subinterfaces" on page 139.
- 3. Create new subinterface(s) and define IP classifiers. This task is optional and only required if you wish to add additional subinterfaces with IP classifiers for further managing traffic from a customer based on the combination of VLAN tags and a specific source IP address, range or subnet.

When creating subinterfaces with VLAN tags and IP classifiers you must first create a subinterface with the same VLAN tag but without any IP classifiers and then create subinterfaces with the IP classifiers. Doing so allows traffic that doesn't match the IP classifier to be assigned to the first interface you created.

You can also use IP classifiers for managing untagged traffic. To do so, you must create a sub-interface with the vlan tag "0", and define sub-interface(s) with IP classifiers for managing untagged traffic using IP classifiers.



Note: IP classification may only be used on the subinterfaces associated with one side of the virtual wire. The subinterfaces defined on the corresponding side of the virtual wire must use the same VLAN tag, but must not include an IP classifier.



Figure 12. Virtual Wire Deployment with Subinterfaces (VLAN Tags only)

The illustration in Figure 12 depicts CustomerA and CustomerB connected to the firewall through one physical interface, ethernet1/1, configured as a Virtual Wire; it is the ingress interface. A second physical interface, ethernet1/2, is also part of the Virtual Wire; it is the egress interface that provides access to the Internet. For CustomerA, you also have subinterfaces ethernet1/1.1 (ingress) and ethernet1/2.1 (egress). For CustomerB, you have the subinterface ethernet1/1.2 (ingress) and

ethernet1/2.2 (egress). When configuring the subinterfaces, you must assign the appropriate VLAN tag and zone in order to apply policies for each customer. In this example, the policies for CustomerA are created between Zone1 and Zone2, and policies for CustomerB are created between Zone3 and Zone4.

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this example, a single subinterface matches the VLAN tag on the incoming packet, hence that subinterface is selected. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.



Note: The same VLAN tag must not be defined on the parent virtual wire interface and the subinterface. Verify that the VLAN tags defined on the **Tag Allowed** list of the parent virtual wire interface (**Network > Virtual Wires**) are not included on a subinterface.

The illustration in Figure 13 depicts CustomerA and CustomerB connected to one physical firewall that has two virtual systems (vsys), in addition to the default virtual system (vsys1). Each virtual system is an independent virtual firewall that is managed separately for each customer. Each vsys has attached interfaces/subinterfaces and security zones that are managed independently.



Figure 13. Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP classifiers)

Vsys1 is set up to use the physical interfaces ethernet1/1 and ethernet1/2 as a virtual wire; ethernet1/1 is the ingress interface and ethernet1/2 is the egress interface that provides access to the Internet. This virtual wire is configured to accept all tagged and untagged traffic with the exception of VLAN tags 100 and 200 that are assigned to the subinterfaces.

CustomerA is managed on vsys2 and CustomerB is managed on vsys3. On vsys2 and vsys3, the following vwire subinterfaces are created with the appropriate VLAN tags and zones to enforce policy measures:

Customer	Vsys	Vwire Subinterfaces	Zone	VLAN Tag	IP Classifier
А	2	e1/1.1 (ingress)	Zone3	100	None
		e1/2.1 (egress)	Zone4	100	
	2	e1/1.2 (ingress)	Zone5	100	IP subnet 192.1.0.0/
		e1/2.2 (egress)	Zone6	100	16
	2	e1/1.3 (ingress)	Zone7	100	IP subnet 192.2.0.0/
	e1/2.3 (egress)	Zone8	100	16	
В	3	e1/1.4 (ingress)	Zone9	200	None
		e1/2.4 (egress)	Zone10	200	

When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this case, for CustomerA, there are multiple subinterfaces that use the same VLAN tag. Hence, the firewall first narrows the classification to a subinterface based on the source IP address in the packet. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

For return-path traffic, the firewall compares the destination IP address as defined in the IP classifier on the customer-facing subinterface and selects the appropriate virtual wire to route traffic through the accurate subinterface.



Note: When a session(s) created on a subinterface with VLAN and IP classifiers is offloaded to the hardware for forwarding, the interface counter on the VLAN interface is incremented instead of the subinterface with the IP classifiers. The session counters, however, are accurate for all the interfaces and subinterfaces.

Layer 2 Deployments

In a Layer 2 deployment, the firewall provides switching between two or more networks. Each group of interfaces must be assigned to a VLAN object in order for the firewall to switch between them. The firewall will perform VLAN tag switching when layer 2 subinterfaces are attached to a common VLAN object. Choose this option when switching is required (Figure 14).



Figure 14. Layer 2 Deployment

Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between multiple ports. An IP address must be assigned to each interface and a virtual router must be defined to route the traffic. Choose this option when routing is required (Figure 15).



Figure 15. Layer 3 Deployment

Point-to-Point Protocol over Ethernet Support

You can configure the firewall to be a Point-to-Point Protocol over Ethernet (PPPoE) termination point to support connectivity in a Digital Subscriber Line (DSL) environment where there is a DSL modem but no other PPPoE device to terminate the connection.

You can choose the PPPoE option and configure the associated settings when an interface is defined as a Layer 3 interface. For instructions, refer to "Configuring Layer 3 Interfaces" on page 130.

Note: PPPoE is not supported in HA active/active mode.

DHCP Client

You can configure the firewall interface to act as a DHCP client and receive a dynamically assigned IP address. The firewall also provides the capability to propagate settings received by the DHCP client interface into a DHCP server operating on the firewall. This is most commonly used to propagate DNS server settings from an Internet service provider to client machines operating on the network protected by the firewall.



Note: DHCP client is not supported in HA active/active mode.

Tap Mode Deployments

A network tap is a device that provides a way to access data flowing across a computer network. Tap mode deployment allows you to passively monitor traffic flows across a network by way of a switch SPAN or mirror port.

The SPAN or mirror port permits the copying of traffic from other ports on the switch. By dedicating an interface on the firewall as a tap mode interface and connecting it with a switch SPAN port, the switch SPAN port provides the firewall with the mirrored traffic. This provides application visibility within the network without being in the flow of network traffic.



Note: When deployed in tap mode, the firewall is not able to take action, such as blocking traffic or applying QoS traffic control.

Defining Virtual Wires

Network > Virtual Wires

Use this page to define virtual wires after you have specified two virtual wire interfaces on the firewall. For an overview of virtual wire deployments, refer to "Virtual Wire Deployments" on page 120. For instructions on specifying interfaces as virtual wire, refer to "Configuring Virtual Wire Interfaces" on page 138.

Field	Description
Virtual Wire Name	Enter a virtual wire name (up to 31 characters). This name appears in the list of virtual wires when configuring interfaces. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Interfaces	Select two Ethernet interfaces from the displayed list for the virtual wire configuration. Interfaces are listed here only if they have the virtual wire interface type and have not been assigned to another virtual wire.

Table 47.	Virtual	Wire	Settings
-----------	---------	------	----------

Field	Description
Tags Allowed	Enter the tag number (0 to 4094) or range of tag numbers (tag1-tag2) for the traffic allowed on the virtual wire. A tag value of zero indicates untagged traffic (the default). Multiple tags or ranges must be separated by commas. Traffic that has an excluded tag value is dropped. Note that tag values are not changed on incoming or outgoing packets.
	When utilizing virtual wire subinterfaces, the Tag Allowed list will cause all traffic with the listed tags to be classified to the parent virtual wire. Virtual wire subinterfaces must utilize tags that do not exist in the parent's Tag Allowed list.
Multicast Firewalling	Select this option if you want to be able to apply security rules to multicast traffic. If this setting is not enabled, multicast traffic is forwarded across the virtual wire.
Link State Pass Through	Select this check box if you want to bring down the other port in a virtual wire when a down link state is detected. If this check box is not selected, link status is not propagated across the virtual wire.

 Table 47.
 Virtual Wire Settings (Continued)

To change a virtual wire name or the allowed tags, click the virtual wire name on the **Virtual Wires** page, change the settings, and click **OK**. Virtual wires also can be changed from the **Interfaces** page (refer to "Configuring Virtual Wire Interfaces" on page 138).

To delete one or more virtual wires, select the check box next to the virtual wire names and click **Delete**. Note that deleting a virtual wire removes it from the associated virtual wire interfaces shown on the **Interfaces** page.

Packet Content Modification

Palo Alto Networks firewalls can be configured to perform the following packet content modifications:

- By default, the firewall will clear TCP urgent flags from all packets. This behavior can be changed by disabling this feature in the CLI using the following command: set deviceconfig setting tcp urgent-data oobinline.
- Modification of the DiffServ Code Point (DSCP) or IP precedences QoS markings can be configured in the **Actions** section of each security rule.
- VLAN tag translation will be performed whenever the firewall's input and output interfaces are tagged differently.

Note: VLAN tag translation is not available on virtual wire interfaces.

Firewall Interfaces

The following table describes the types of interfaces supported on the firewall and how to define them.

Interface	Description	
Layer 2	One or more Layer 2 interfaces can be configured for untagged VLAN traffic. You can then define Layer 2 subinterfaces for traffic with specific VLAN tags. Refer to "Configuring Layer 2 Interfaces" on page 86 and "Configuring Layer 2 Subinterfaces" on page 87.	
Layer 3	One or more Layer 3 interfaces can be configured for untagged routed traffic. You can then define Layer 3 subinterfaces for traffic with specific VLAN tags.	
	Each interface can have multiple IP addresses. Refer to "Configuring Layer 3 Interfaces" on page 88 and "Configuring Layer 3 Subinterfaces" on page 91.	
Aggregate Ethernet	Two or more Ethernet ports can be combined into a group to increase the throughput and resiliency for a Layer 2, Layer 3, or virtual wire interface and its subinterfaces. Refer to "Configuring Aggregate Ethernet Interfaces" on page 142.	
	<i>Note:</i> You cannot apply QoS settings to an aggregate Ethernet interface.	
VLAN	VLAN interfaces provide Layer 3 routing of VLAN traffic to non-VLAN destinations. Refer to "Configuring VLAN Interfaces" on page 143.	
Loopback	Loopback interfaces can be used to provide Layer 3 services such as in- band management, GlobalProtect portal or gateway functionality, and IPSec. Each loopback interface behaves as a host interface and is assigned an IP address. Refer to "Configuring Loopback Interfaces" on page 146.	
Tunnel	Tunnel interfaces can be configured. Refer to "Configuring Tunnel Interfaces" on page 147.	
Virtual Wire	A virtual wire binds two Ethernet ports together, which allows you to install the firewall transparently in the network with minimum configuration. A virtual wire accepts all traffic or traffic with selected VLAN tags, but provides no switching or routing services. You can also create subinterfaces and classify traffic according to an IP address, IP range, or subnet. Refer to "Configuring Virtual Wire Interfaces" on page 138.	
Тар	The Tap interface permits connection to a span port on a switch for traffic monitoring only. This mode does not support traffic blocking or URL filtering. Refer to "Configuring Tap Interfaces" on page 149.	
High Availability	You can configure a data interface to be a high availability (HA) interface on some Palo Alto Networks firewalls. Refer to "Configuring HA Interfaces" on page 150.	

 Table 48.
 Supported Interfaces

Viewing the Current Interfaces

► Network > Interfaces

The **Interfaces** page lists the interface type, link state, and security zone for each configured interface, along with the IP address, virtual router, VLAN tag, and VLAN or virtual wire name (as applicable).

By default, the interfaces are listed by interface name.

The following icon is used on the Interfaces page:

Indicates the link is up (green), down (red), or in an unknown state (gray).

Configuring Layer 2 Interfaces

Network > Interfaces > Ethernet

You can configure one or more Ethernet ports as Layer 2 interfaces for untagged VLAN traffic. For each main Layer 2 interface, you can define multiple Layer 2 subinterfaces for traffic with specific VLAN tags (refer to "Configuring Layer 2 Subinterfaces" on page 129). VLAN interfaces can also be used in Layer 2 deployments to provide routing or gateway services to traffic in the Layer 2 domain (refer to "Configuring VLAN Interfaces" on page 143). This can be achieved by creating a VLAN interface and assigning it as the default gateway for hosts connected to the firewall's Layer2 interface.

To configure a Layer 2 Ethernet interface, click the link for the interface on the **Ethernet** tab, and specify the following settings.

Field	Description	
Interface Name	Choose the interface from the drop-down list. Modify the name if desired.	
Interface Type	Select Layer 2 from the drop-down list.	
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.	
Comment	Enter an optional description of the interface.	
Config Tab		
VLAN	Select a VLAN, or click New to define a new VLAN (refer to "VLAN Support" on page 152). None removes the configuration from the interface. A VLAN object must be configured to enable switching between Layer 2 interfaces or to enable routing through a VLAN interface.	
Virtual System	Select the virtual system for the interface. None removes the configuration from the interface.	
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151). None removes the configuration from the interface.	
Advanced Tab		
Link Speed	Select the interface speed in Mbps (10, 100, or 1000) or select auto.	

Table 49. Layer 2 Interface Settings

Field	Description
Link Duplex	Select whether the interface transmission mode is full-duplex (full), half- duplex (half), or negotiated automatically (auto).
Link State	Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto).

 Table 49.
 Layer 2 Interface Settings (Continued)

Configuring Layer 2 Subinterfaces

► Network > Interfaces

For each Ethernet port configured as a Layer 2 interface, you can define an additional logical Layer 2 interface (subinterface) for each VLAN tag that is used on the traffic received by the port. To configure the main Layer 2 interfaces, refer to "Configuring Layer 2 Interfaces" on page 128. To enable switching between Layer 2 subinterfaces, simply link those subinterfaces to the same VLAN object.

To add a Layer 2 Ethernet subinterface, select the associated physical interface and then click **Add Subinterface** and specify the following information.

Field	Description
Interface Name	Select the Layer 2 interface where you want to add a subinterface. To configure the Layer 2 interfaces, refer to "Configuring Layer 2 Interfaces" on page 128.
	Enter the number (1 to 9999) appended to the physical interface name to form the logical interface name. The general name format is:
	ethernet <i>x</i> / <i>y</i> .<1-9999>
Tag	Enter the tag number (1 to 4094) of the traffic received on this interface. Outgoing traffic on this interface is also set to this tag value.
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Enter an optional description of the interface.
Assign Interface To	
VLAN	For a Layer 2 interface, select a VLAN, or click New to define a new VLAN (refer to "Network Profiles" on page 174). None removes the configuration from the interface. A VLAN object must be configured to enable switching between Layer 2 interfaces or to enable routing through a VLAN interface.
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151). None removes the configuration from the interface.
Virtual System	Select the virtual system for the interface. None removes the configuration from the interface.

Table 50. Layer 2 Subinterface Settings

Configuring Layer 3 Interfaces

▶ Network > Interfaces > Ethernet

You can configure one or more Ethernet ports as Layer 3 interfaces for untagged routed traffic. You can then define Layer 3 subinterfaces for traffic with specific VLAN tags (refer to "Configuring Layer 3 Subinterfaces" on page 134) when connecting the firewall to a neighboring device using a trunk link. For information on configuring Layer 3 interfaces for PPPoE, refer to "Point-to-Point Protocol over Ethernet Support" on page 124. Layer 3 interfaces can also act as a DHCP client to receive its configuration from an external DHCP server.

To configure a Layer 3 Ethernet interface, click the link for the interface on the **Ethernet** tab, and specify the following settings.

Field	Description
Interface Name	Choose the interface from the drop-down list. Modify the name if desired.
Interface Type	Select Layer 3 from the drop-down list.
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Enter an optional description of the interface.
Config Tab	
Virtual Router	Select a virtual router, or click New to define a new virtual router (refer to "Virtual Routers and Routing Protocols" on page 153). None removes the configuration from the interface.
Virtual System	Select the virtual system for the interface. None removes the configuration from the interface.
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151). None removes the configuration from the interface.
IPv4 Tab	
Туре	Choose how the IP address information will be specified (Static, PPPoE, or DHCP Client), as described below.
Static	Enter an IP address and network mask for the interface in the format <i>ip_address/mask</i> , and click Add . You can enter multiple IP addresses for the interface. To delete an IP address, select the address and click Delete .

Table 51. Layer 3 Interface Settings

Field	Description
РРРоЕ	Choose PPPoE if the interface will be used for PPPoE and configure the following settings:
	General subtab:
	• Enable —Select the check box to activate the interface for PPPoE termi- nation.
	• Username—Enter the user name for the point-to-point connection.
	• Password/Confirm Password —Enter and then confirm the password for the user name.
	Advanced subtab:
	• Authentication—Choose CHAP (Challenge-Handshake Authentication Protocol), PAP (Password Authentication Protocol), or the default Auto (to have the firewall determine the authentication protocol for PPPoE communications).
	• Static Address—Specify the static IP address that was assigned by the service provider (optional, no default).
	• Automatically create default route pointing to peer—Select the check box to automatically create a default route that points to the PPPoE peer when connected.
	• Default Route Metric —Specify the route metric to be associated with the default route and used for path selection (optional, range 1-65535).
	• Access Concentrator—Specify the name of the access concentrator to which the connection is made (optional, no default).
	• Service—Specify the service string (optional, no default).
	• Passive —Select the check box to use passive mode. In passive mode, a PPPoE end point waits for the access concentrator to send the first frame.
	<i>Note:</i> PPPoE is not supported in HA active/active mode.
DHCP Client	Choose DHCP Client to allow the interface to act as a DHCP client and receive a dynamically assigned IP address. Specify the following:
	• Enable —Select the check box to activate the DHCP client on the interface.
	• Automatically create default route point to server—Select the check box to automatically create a default route that points to the default gateway provided by the DHCP server.
	• Default Route Metric —Specify the route metric to be associated with the default route and used for path selection (optional, range 1-65535).
	Click Show DHCP Client Runtime Info to open a window that displays all settings received from the DHCP server, including DHCP lease status, dynamic IP assignment, subnet mask, gateway, server settings (DNS, NTP, domain, WINS, NIS, POP3, and SMTP).
	<i>Note:</i> DHCP client is not supported in HA active/active mode.
IPv6 Tab	

Table 51. Layer 3 Interface Settings (Continued)

Enable IPv6 on the	Select the check box to enable IPv6 addressing on this interface.
interface	-

Field	Description
Interface ID	Enter the 64-bit extended unique identifier in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. If the interface ID is left blank, the firewall will use the EUI-64 generated from the physical interface's MAC address. The interface ID is used as the host portion of an interface address when the Use interface ID as host portion option is enabled.
Address	Click Add and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. Select Use interface ID as host portion to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select Anycast to include routing through the nearest node. If the Prefix is not entered, the IPv6 address assigned to the interface will be wholly specified in the address text box.
	Use the Send Router Advertisement (Send RA) option to enable router advertisement for this IP address. You can also set the Autonomous flag to be sent and you can set the on-link option. You must enable the global Enable Router Advertisement option on the interface before enabling Send Router Advertisement option for a specific IP address.
Address Resolution	Select the check box to enable Duplicate Address Detection (DAD) and specify the following information.
(Detection)	• DAD Attempts —Specify the number of attempts within the neighbor solicitation interval for DAD before the attempt to identify neighbors fails (range 1-10).
	• Reachable Time —Specify the length of time that a neighbor remains reachable after a successful query and response (range 1-36000 seconds).
	• Neighbor Solicitation (NS) Interval—Specify the number of seconds for DAD attempts before failure is indicated (range 1-10 seconds).

Table 51. Layer 3 Interface Settings (Continued)

Field	Description
Enable Router Advertisement	Select the check box to enable Router Advertisement (RA) to provide Stateless Address Autoconfiguration (SLAAC) on IPv6 interfaces. This enables the firewall to act as a default gateway for IPv6 hosts that are not statically configured and will provide the host with an IPv6 prefix that can be used for address configuration. A separate DHCPv6 server can be used in conjunction with this feature to provide DNS and other settings to clients.
	This option is a global setting for the interface, you can also set router advertisement options per IP address by clicking Add and entering in an IP address. You must enable this option on the interface if you are going to specify the Send Router Advertisement option per address.
	Specify the following information that will be used by clients who receive the RA messages.
	• Min Interval (sec) —Specify the minimum interval per second that the firewall will send out router advertisements. Router advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 3-1350 seconds, default 200 seconds).
	• Max Interval (sec)—Specify the maximum interval per second that the firewall will send out router advertisements. Router advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 4-1800 seconds, default 600 seconds).
	• Hop Limit —Specify the hop limit that will be applied to clients for out- going packets. Enter 0 for no hop limit (range 1-255, default 64).
	 Link MTU—Specify the link MTU that will be applied to clients. Select unspecified for no link MTU (range 1280-9192, default unspecified).
	• Reachable Time (ms) —Specify the reachable time that the client will use to assume a neighbor is reachable after having received a reachability confirmation message. Select unspecified for no reachable time value (range 0-3600000 milliseconds, default unspecified).
	• Retrans Time (ms) —Specify the retransmission timer that the client will use to determine how long it should wait before retransmitting neighbor solicitation messages. Select unspecified for no retrans time (range 0-4294967295 milliseconds, default unspecified).
	• Router Lifetime (sec) —Specify the router lifetime that instructs the client on how long the firewall/router should be used as the default router (range 0-9000 seconds, default 1800).
	• Managed Configuration—Select the check box to indicate to the client that addresses are available via DHCPv6.
	• Other Configuration—Select the check box to indicate to the client that other addresses information is available via DHCPv6, such as DNS-related settings.
	• Consistency check —Select the check box to enable consistency checks that the firewall will use to verify that router advertisement sent from other routers are advertising consistent information on the link. If inconsistencies are detected, a log will be created.
Advanced Tab	
Link Speed	Select the interface speed in Mbps (10, 100, or 1000) or select auto.
Link Duplex	Select whether the interface transmission mode is full-duplex (full), half- duplex (half), or negotiated automatically (auto).

 Table 51.
 Layer 3 Interface Settings (Continued)

Field	Description
Link State	Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto).
Other Info	Specify the following information on the Other Info subtab:
	• Management Profile—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.
	• MTU—Enter the maximum transmission unit (MTU) in bytes for packets sent on this Layer 3 interface (512 to 1500, default 1500). If machines on either side of the firewall perform Path MTU Discovery (PMTUD), the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.
	• Adjust TCP MSS—If you select this check box, the maximum segment size (MSS) is adjusted to 40 bytes less than the interface MTU. This setting addresses the situation in which a tunnel through the network requires a smaller MSS. If a packet cannot fit within the MSS without fragmenting, this setting allows an adjustment to be made.
	 Untagged Subinterface—Specifies that all subinterfaces belonging to this Layer 3 interface are untagged. PAN-OS selects an untagged subinterface as the ingress interface based on the destination of the packet. If a packet has an untagged subinterface's IP address as the destination, it will map to the subinterface. This also means that packets going in the reverse direction must have their source address translated to the untagged subinterface's interface IP address. A by product of this classification mechanism is that all multicast and broadcast packets will be assigned to the base interface rather than any of the subinterfaces. Since OSPF uses multicast, it is not supported on untagged subinterfaces.
ARP/Interface Entries	To add one or more static ARP entries, click Add and enter an IP address and its associated hardware (MAC) address and Layer 3 interface that can access the hardware address.
ND Entries	Click Add to enter the IPv6 address and MAC address of neighbors to add for discovery.

Table 51. Layer 3 Interface Settings (Continued)

Configuring Layer 3 Subinterfaces

► Network > Interfaces

For each Ethernet port configured as a Layer 3 interface, you can define an additional logical Layer 3 interface (subinterface) for each VLAN tag that is used on the traffic received by the port. To configure the main Layer 3 interfaces, refer to "Configuring Layer 3 Interfaces" on page 130.

Untagged layer 3 subinterfaces may also be used when the parent Layer 3 interface's "untagged subinterface" option is enabled. Untagged subinterfaces are used in multi-tenant environments where each tenant's traffic must leave the firewall without VLAN tags.

Consider an example where each tenant's traffic egresses the firewall and the next hop is an ISP router. It is not always possible to apply a VLAN tag on the return traffic for proper classification into a virtual system by the firewall. In these cases, you can use an untagged subinterface on the ISP-router facing side. Each untagged subinterface will have an IP address and all outgoing traffic must have its source address translated to that interface IP address. An explicit NAT rule must be created for this feature to function. Source NAT is required on the untagged subinterfaces because the firewall will use the destination IP address on inbound

(return path) packets to select the appropriate virtual system for policy lookup. Any traffic received on the parent interface that is not destined for one of the untagged subinterface IPs will be handled by the virtual system and virtual router assigned to that parent interface.

To add a Layer 3 Ethernet subinterface, select the associated interface and click **Add Subinterface** and specify the following information.

Field	Description
Interface Name	Select the Layer 3 interface where you want to add a subinterface. To configure the Layer 3 interfaces, refer to "Configuring Layer 3 Interfaces" on page 130.
	Enter the number (1 to 9999) appended to the physical interface name to form the logical interface name. The general name format is:
	ethernet <i>x</i> / <i>y</i> .<1-9999>
Tag	Enter the tag number (1 to 4094) of the traffic received on this interface. Outgoing traffic on this interface is also set to this tag value.
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Enter an optional description of the interface.
Config Tab	
Virtual Router	Select a virtual router, or click New to define a new virtual router (refer to "Virtual Routers and Routing Protocols" on page 153). None removes the configuration from the interface.
Virtual System	Select the virtual system for the interface. None removes the configuration from the interface.
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151). None removes the configuration from the interface.
IPv4 Tab	
Туре	Choose how the IP address information will be specified (Static, PPPoE, or DHCP Client), as described below.
Static	Enter an IP address and network mask for the interface in the format <i>ip_address/mask</i> , and click Add . You can enter multiple IP addresses for the interface. To delete an IP address, select the address and click Delete .
DHCP Client	Choose DHCP Client to allow the interface to act as a DHCP client and receive a dynamically assigned IP address. Specify the following:
	• Enable —Select the check box to activate the DHCP client on the inter- face.
	• Automatically create default route pointing to default gateway pro- vided by server—Select the check box to automatically create a default route that points to the DHCP server when connected.
	• Default Route Metric —Specify the route metric to be associated with the default route and used for path selection (optional, range 1-65535).
	• Show DHCP Client Runtime Info—Select to open a window that displays all settings received from the DHCP server, including DHCP lease status, dynamic IP assignment, subnet mask, gateway, server settings (DNS, NTP, domain, WINS, NIS, POP3, and SMTP).

 Table 52.
 Layer 3 Subinterface Settings

Field	Description
IPv6 Tab	
Enable IPv6 on the interface	Select the check box to enable IPv6 addressing on this interface.
Interface ID	Enter the 64-bit extended unique identifier in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. If the interface ID is left blank, the firewall will use the EUI-64 generated from the physical interface's MAC address.
Address	Click Add and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. Select Use interface ID as host portion to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select Anycast to include routing through the nearest node. If Prefix is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.
	Use the Send Router Advertisement (Send RA) option to enable router advertisement for this IP address. You can also set the Autonomous flag to be sent and you can set the on-link option. You must enable the global Enable Router Advertisement option on the interface before enabling Send RA option for a specific IP address.
Address Resolution (Duplicate Address	Select the check box to enable Duplicate Address Detection (DAD) and specify the following information.
(Detection)	• DAD Attempts —Specify the number of attempts within the neighbor solicitation interval for DAD before the attempt to identify neighbors fails (range 1-10).
	• Reachable Time —Specify the length of time that a neighbor remains reachable after a successful query and response (range 1-36000 seconds).
	• Neighbor Solicitation (NS) Interval —Specify the number of seconds for DAD attempts before failure is indicated (range 1-10 seconds).

Table 52. Layer 3 Subinterface Settings (Continued)

Field	Description
Enable Router Advertisement	Select the check box to enable Router Advertisement (RA) to provide Stateless Address Autoconfiguration (SLAAC) on IPv6 interfaces. This enables the firewall to act as a default gateway for IPv6 hosts that are not statically configured and will provide the host with an IPv6 prefix that can be used for address configuration. A separate DHCPv6 server can be used in conjunction with this feature to provide DNS and other settings to clients.
	This option is a global setting for the interface, you can also set router advertisement options per IP address by clicking Add and entering in an IP address. You must enable this option on the interface if you are going to specify the Send Router Advertisement option per address.
	Specify the following information that will be used by clients who receive the RA messages.
	• Min Interval (sec) —Specify the minimum interval per second that the firewall will send out router advertisements. Router Advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 3-1350 seconds, default 200 seconds).
	• Max Interval (sec)—Specify the maximum interval per second that the firewall will send out router advertisements. Router Advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 4-1800 seconds, default 600 seconds).
	• Hop Limit—Specify the hop limit that will be applied to clients for out- going packets. Enter 0 for no hop limit (range 1-255, default 64).
	• Link MTU—Specify the link MTU that will be applied to clients. Select unspecified for no link MTU (range 1280-9192, default unspecified).
	• Reachable Time (ms) —Specify the reachable time that the client will use to assume a neighbor is reachable after having received a reachability confirmation message. Select unspecified for no reachable time value (range 0-3600000 milliseconds, default unspecified).
	• Retrans Time (ms) —Specify the retransmission timer that the client will use to determine how long it should wait before retransmitting neighbor solicitation messages. Select unspecified for no retrans time (range 0-4294967295 milliseconds, default unspecified).
	• Router Lifetime (sec) —Specify the router lifetime that instructs the client on how long the firewall/router should be used as the default router (range 0-9000 seconds, default 1800).
	• Managed Configuration—Select the check box to indicate to the client that addresses are available via DHCPv6.
	• Other Configuration—Select the check box to indicate to the client that other addresses information is available via DHCPv6, such as DNS-related settings.
	• Consistency check —Select the check box to enable consistency checks that the firewall will use to verify that router advertisement sent from other routers are advertising consistent information on the link. If inconsistencies are detected, a log will be created.

Table 52. Layer 3 Subinterface Settings (Continued)

Field	Description
Advanced Tab	
Other Info	Specify the following information on the Other Info subtab:
	• Management Profile—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.
	• MTU—Enter the maximum transmission unit (MTU) in bytes for packets sent on this Layer 3 interface (512 to 1500, default 1500). If machines on either side of the firewall perform Path MTU Discovery (PMTUD), the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.
	• Adjust TCP MSS—If you select this check box, the maximum segment size (MSS) is adjusted to 40 bytes less than the interface MTU. This setting addresses the situation in which a tunnel through the network requires a smaller MSS. If a packet cannot fit within the MSS without fragmenting, this setting allows an adjustment to be made.
ARP Entries	To add one or more static Address Resolution Protocol (ARP) entries, enter an IP address and its associated hardware (Media Access Control or MAC) address, and click Add . To delete a static entry, select the entry and click Delete . Static ARP entries reduce ARP processing and preclude man-in-the-middle attacks for the specified addresses.
ND Entries	Click Add to enter the IP address and MAC address of neighbors to add for discovery.

Table 52. Layer 3 Subinterface Settings (Continued)

Configuring Virtual Wire Interfaces

► Network > Interfaces

To create a virtual wire, you bind two Ethernet ports together, which allows all traffic to pass between the ports, or just traffic with selected VLAN tags (no other switching or routing services are available). You can also create subinterfaces and classify traffic according to an IP address, IP range, or subnet. A virtual wire requires no changes to adjacent network devices. For an overview of virtual wire deployments, refer to "Virtual Wire Deployments" on page 120.

To set up a virtual wire through the firewall, you must first define the virtual wire interfaces, as described in the following procedure. You then create the virtual wire using the interfaces that you created.

To configure each virtual wire interface, follow these steps:

- 1. Identify the interface you want to use for the virtual wire on the **Ethernet** tab, and remove it from the current security zone, if any.
- 2. Click the interface name and specify the following information.

Field	Description
Interface Name	The interface name is automatically populated based on the interface that you selected. This field cannot be edited.
Interface Type	Select Virtual Wire from the drop-down list.

 Table 53.
 Virtual Wire Settings

Field	Description
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Enter an optional description of the interface.
Config Tab	
Virtual Wire	Select a virtual wire, or click New to define a new virtual wire (refer to "Defining Virtual Wires" on page 125).
Virtual System (only on systems with multi- virtual system capability)	Select the virtual system for the interface.
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151).
Advanced Tab	
Link Speed	Specify the interface speed. If the selected interface is a 10 Gbps interface, the only option is auto . In other cases, the options are: 10 , 100 , 1000 , or auto .
Link Duplex	Select whether the interface transmission mode is full-duplex (full), half- duplex (half), or negotiated automatically (auto).
Link State	Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto).

Table 53. Virtual Wire Settings (Continued)

If you want to change a virtual wire to another interface type, click the virtual wire name shown in the **VLAN/Virtual Wire** column, if any, select **None**, and click **OK**.

Configuring Virtual Wire Subinterfaces

► Network > Interfaces

Using virtual wire subinterfaces allows you to separate traffic by VLAN tags or a VLAN tag and IP classifier combination, assign the tagged traffic to a different zone and virtual system, and then enforce security policies for the traffic that matches the defined criteria.

For each Ethernet port configured as a virtual wire interface, you can create one or more logical subinterfaces. Each subinterface can segregate traffic using a VLAN tag (with a zero or non-zero value, where the value zero indicates untagged traffic). After you set up a VLAN tagged subinterface, you can then create another subinterface that segregates traffic by combining the VLAN tag with IP classifiers such as a single IP address, a range, or a subnet.

IP classification on subinterfaces can only be used on one side of a virtual wire. Traffic entering the firewall through a virtual wire subinterface where IP classification is defined will be classified by its source IP address. For return path traffic entering through the other end of the virtual wire, classification is based on the destination IP address.

Make sure that the VLAN tags defined on the **Tag Allowed** list of the parent virtual wire interface (Network > Virtual Wires) are not included on the list of tags defined on a virtual wire subinterface.

To configure a parent or main virtual wire interfaces, refer to "Configuring Virtual Wire Interfaces" on page 138.



Note: The following NAT limitations apply to virtual wires:

- Source NAT where the ingress interface uses source IP based classification is not supported.
- Destination NAT cannot be used where the egress interface uses source IP based classification.

To add a virtual wire subinterface, select the virtual wire interface where you want to add the subinterface, and click **Add Subinterface** and specify the following information.

Field	Description
Interface Name	The main interface name is automatically populated based on the interface that you selected; the label cannot be edited.
	To define the subinterface, enter a number (1 to 9999) to the physical interface name to form the logical interface name. The general name format is:
	ethernet <i>x</i> / <i>y</i> .<1-9999>
	To configure the virtual wire, refer to "Configuring Virtual Wire Interfaces" on page 138.
Tag	Enter the tag number (0 to 4094) of the traffic received on this interface.
	Setting a tag value of 0 will match untagged traffic.
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Enter an optional description of the interface.
IP Classifier	Click Add to add an IP address, subnet, or IP range or any combination of the IP classifiers, to classify traffic entering the firewall through this physical port into this subinterface based on its source IP address. Return- path traffic entering the firewall through the other end of the associated virtual wire will be matched according to its destination address.
	On a virtual wire subinterface, IP classification can only be used in conjunction with VLAN based classification.
Assign Interface To	
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151).
Virtual System (only on systems with multi- virtual system capability)	Select the virtual system for the interface.
Virtual Wire	Select a virtual wire, or click New to define a new virtual wire (refer to "Defining Virtual Wires" on page 125).

Table 54. Virtual Wire Subinterface Settings

Configuring Aggregate Interface Groups

► Network > Interfaces

Aggregate interface groups allow you to generate more than 1 Gbps aggregate throughput by using 802.3ad link aggregation of multiple 1 Gbps links. Aggregation of 10Gbps XFP and SFP+ is also supported. The aggregate interface that you create becomes a logical interface. Interface management, zone profiles, VPN interfaces, and VLAN subinterfaces are all properties of the logical aggregate interface, not of the underlying physical interfaces.

Each aggregate group can contain several physical interfaces of the type Aggregate Ethernet. After the group is created, you perform operations such as configuring Layer 2 or Layer 3 parameters on the Aggregate Group object rather than on the Aggregate Ethernet interfaces themselves.



Note: The algorithm used to evenly distribute traffic on the interfaces in an aggregate group is based on the session ID of the incoming traffic. As traffic enters the aggregate group, the last three unique bits of the session ID are used to determine which interface should be used for a given flow.

The following rules apply to aggregate interface groups:

- The interfaces are compatible with virtual wire, Layer 2, and Layer 3 interfaces.
- Tap mode is not supported.
- The 1 Gbps links in a group must be of the same type (all copper or all fiber).
- You can include up to eight aggregate interfaces in an aggregate group.
- All of the members of an aggregate group must be of the same type. This is validated during the commit operation.
- Aggregate groups can be used for redundancy and throughput scaling on the HA3 (packet forwarding) link in Active/Active HA deployments.

You can configure one or more interfaces as part of an aggregate Ethernet interface group. First define the group, as described in this section, and then assign interfaces to the group. For instructions on assigning interfaces to the group, refer to "Configuring Layer 3 Subinterfaces" on page 134.

To create and configure aggregate group interfaces, click **Add Aggregate Group** and specify the following information.

Field	Description
Interface Name	Enter a name and numeric suffix to identify the interface. The interface name is listed as $mm.n$ where mm is the name and n is the suffix (1-8).
Interface Type	Select the interface type.
	• HA—No additional configuration is required.
	• Layer 2—Configure the settings as described in Table 50.
	• Layer 3—Configure the settings as described in Table 52.
Comment	Enter an optional description of the interface.

Table 55. Aggregate Group Interface Settings

Field	Description
Assign Interface To	
Assign Interface To	The interface assignment depends on the interface type, as follows:
	• Layer 2—Specify a VLAN and zone.
	• Layer 3—Specify a virtual router and zone
	• Virtual Wire—Specify a virtual wire and zone.
	<i>Note:</i> If the type is HA, there are no options to specify in this section.
Virtual System	Select the virtual system for the interface. None removes the configuration from the interface.

Table 55. Aggregate Group Interface Settings (Continued)

Configuring Aggregate Ethernet Interfaces

► Network > Interfaces

Each aggregate Ethernet interface is assigned a name of the form ae.*number* and can be of the type Layer 2, Layer 3, or virtual wire. After the assignment is made, the new interface functions in the same way as any other interface.

To configure aggregate Ethernet interfaces, click the interface name on the **Ethernet** tab and specify the following information.

Field	Description
Interface Name	Choose the interface from the drop-down list. Modify the name if desired.
Interface Type	Select Aggregate Ethernet from the drop-down list.
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Enter an optional description of the interface.
Config Tab	
Virtual System	Select the virtual system for the interface. None removes the configuration from the interface.
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151). None removes the configuration from the interface.
Advanced Tab	
Link Speed	Select the interface speed in Mbps (10, 100, or 1000) or auto.
Link Duplex	Select whether the interface transmission mode is full-duplex (full), half- duplex (half), or negotiated automatically (auto).
Link State	Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto).

Table 56. Aggregate Ethernet Interface Settings

Configuring VLAN Interfaces

► Network > Interfaces

For each Ethernet port configured as a Layer 2 interface, you can define a VLAN interface to allow routing of the VLAN traffic to Layer 3 destinations outside the VLAN. To configure the main Layer 2 interfaces, refer to "Configuring Layer 2 Interfaces" on page 128.

To define a VLAN interface, open the **VLAN** tab, click **Add**, and specify the following settings.

Field	Description
Interface Name	Specify a numeric suffix for the interface (1-4999).
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Add an optional description of the interface.
Config Tab	
VLAN	Select a VLAN, or click New to define a new VLAN (refer to "Network Profiles" on page 174). None removes the configuration from the interface.
Virtual Router	Select a virtual router, or click New to define a new virtual router (refer to "Virtual Routers and Routing Protocols" on page 153). None removes the configuration from the interface.
Virtual System	Select the virtual system for the interface. None removes the configuration from the interface.
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151). None removes the configuration from the interface.
IPv4 Tab	
Static	Select Static to assign static IP addresses. Click Add and enter an IP address and network mask for the interface in the format <i>ip_address/mask</i> . You can enter multiple IP addresses for the interface.
DHCP Client	Select DHCP to use DHCP address assignment for the interface, and specify the following:
	• Enable —Select the check box to activate the DHCP client on the inter- face.
	• Automatically create default route point to server—Select the check box to automatically create a default route that points to the DHCP server when connected.
	• Default Route Metric —Specify the route metric to be associated with the default route and used for path selection (optional, range 1-65535).
	Click Show DHCP Client Runtime Info to open a window that displays all settings received from the DHCP server, including DHCP lease status, dynamic IP assignment, subnet mask, gateway, server settings (DNS, NTP, domain, WINS, NIS, POP3, and SMTP).

Table 57. VLAN Interface Settings

Field	Description
ARP Entries	To add one or more static ARP entries, enter an IP address and its associated hardware (MAC) address, and click Add . To delete a static entry, select the entry and click Delete .
IPv6 Tab	
Enable IPv6 on the interface	Select the check box to enable IPv6 addressing for the subinterface.
Interface ID	Enter the 64-bit extended unique identifier in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. If the interface ID is left blank, the firewall will use the EUI-64 generated from the physical interface's MAC address.
Address	Click Add and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. Select Use interface ID as host portion to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select Anycast to include routing through the nearest node. If Prefix is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.
	Use the Send Router Advertisement (Send RA) option to enable router advertisement for this IP address. You can also set the Autonomous flag to be sent and you can set the on-link option. You must enable the global Enable Router Advertisement option on the interface before enabling Send Router Advertisement option for a specific IP address.
Address Resolution (Duplicate Address Detection	Select the check box to enable Duplicate Address Detection (DAD) and specify the following information.
	• DAD Attempts —Specify the number of attempts within the neighbor solicitation interval for DAD before the attempt to identify neighbors fails (range 1-10).
	• Reachable Time —Specify the length of time that a neighbor remains reachable after a successful query and response (range 1-36000 seconds).
	Neighbor Solicitation (NS) Interval —Specify the number of seconds for DAD attempts before failure is indicated (range 1-10 seconds).

Table 57. VLAN Interface Settings (Continued)
Field	Description
Enable Router Advertisement	Select the check box to enable Router Advertisement (RA) to provide Stateless Address Autoconfiguration (SLAAC) on IPv6 interfaces. This enables the firewall to act as a default gateway for IPv6 hosts that are not statically configured and will provide the host with an IPv6 prefix that can be used for address configuration. A separate DHCPv6 server can be used in conjunction with this feature to provide DNS and other settings to clients.
	This option is a global setting for the interface, you can also set router advertisement options per IP address by clicking Add and entering in an IP address. You must enable this option on the interface if you are going to specify the Send Router Advertisement option per address.
	Specify the following information that will be used by clients who receive the RA messages.
	• Min Interval (sec) —Specify the minimum interval per second that the firewall will send out router advertisements. Router Advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 3-1350 seconds, default 200 seconds).
	• Max Interval (sec)—Specify the maximum interval per second that the firewall will send out router advertisements. Router Advertisements will be sent at random intervals between the minimum and maximum values that are configured (range 4-1800 seconds, default 600 seconds).
	• Hop Limit —Specify the hop limit that will be applied to clients for out- going packets. Enter 0 for no hop limit (range 1-255, default 64).
	• Link MTU—Specify the link MTU that will be applied to clients. Select unspecified for no link MTU (range 1280-9192, default unspecified).
	• Reachable Time (ms) —Specify the reachable time that the client will use to assume a neighbor is reachable after having received a reachability confirmation message. Select unspecified for no reachable time value (range 0-3600000 milliseconds, default unspecified).
	• Retrans Time (ms) —Specify the retransmission timer that the client will use to determine how long it should wait before retransmitting neighbor solicitation messages. Select unspecified for no retrans time (range 0-4294967295 milliseconds, default unspecified).
	• Router Lifetime (sec) —Specify the router lifetime that instructs the client on how long the firewall/router should be used as the default router (range 0-9000 seconds, default 1800).
	• Managed Configuration—Select the check box to indicate to the client that addresses are available via DHCPv6.
	• Other Configuration—Select the check box to indicate to the client that other addresses information is available via DHCPv6, such as DNS-related settings.
	• Consistency check —Select the check box to enable consistency checks that the firewall will use to verify that router advertisement sent from other routers are advertising consistent information on the link. If inconsistencies are detected, a log will be created.

Table 57. VLAN Interface Settings (Continued)

Field	Description
Advanced Tab	
Other Info	Specify the following:
	• Management Profile—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.
	• MTU—Enter the MTU in bytes for packets sent on this interface (512- 1500, default 1500). If machines on either side of the firewall perform PMTUD, the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.
	• Adjust TCP MSS—if you select this check box, the maximum segment size (MSS) is adjusted to 40 bytes less than the interface MTU. This setting addresses the situation in which a tunnel through the network requires a smaller MSS. If a packet cannot fit within the MSS without fragmenting, this setting allows an adjustment to be made.
ARP/Interface Entries	To add one or more static ARP entries, click Add and enter an IP address and its associated hardware (MAC) address and Layer 3 interface that can access the hardware address.
ND Entries	Click Add to enter the IP address and MAC address of neighbors to add for discovery.

Table 57. VLAN Interface Settings (Continued)

Configuring Loopback Interfaces

► Network > Interfaces

You can define one or more Layer 3 loopback interfaces, as needed. For example, you can define a loopback interface to manage the firewall instead of using the management port. To define a loopback interface, open the **Loopback** tab, click **Add**, and specify the following settings.

Field	Description
Interface Name	Specify a numeric suffix for the interface (1-4999).
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Add an optional description of the interface.
Config Tab	
Virtual Router	Select a virtual router, or click New to define a new virtual router (refer to "Virtual Routers and Routing Protocols" on page 153). None removes the configuration from the interface.
Virtual System	Select the virtual system for the interface. None removes the configuration from the interface.
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151). None removes the configuration from the interface.

Table 58. Loopback Interface Settings

Field	Description
IPv4 Tab	
IP Address	Click Add to enter IP addresses and network masks for the interface.
IPv6 Tab	
Enable IPv6 on the interface	Select the check box to enable IPv6 addressing for the subinterface.
Interface ID	Specify the unique 64-bit hexadecimal identifier for the subinterface.
Address	Enter the IPv6 address. Select Use interface ID as host portion to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select Anycast to include routing through the nearest node.
Advanced Tab	
Other Info	Specify the following settings:
	• Management Profile—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.
	• MTU—Enter the maximum transmission unit (MTU) in bytes for packets sent on this interface (512 to 1500, default 1500). If machines on either side of the firewall perform Path MTU Discovery (PMTUD), the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.
	• Adjust TCP MSS—If you select this check box, the maximum segment size (MSS) is adjusted to 40 bytes less than the interface MTU. This setting addresses the situation in which a tunnel through the network requires a smaller MSS. If a packet cannot fit within the MSS without fragmenting, this setting allows an adjustment to be made.

 Table 58.
 Loopback Interface Settings (Continued)

Configuring Tunnel Interfaces

► Network > Interfaces

To define tunnel interfaces, open the **Tunnel** tab, click **Add**, and specify the following settings.

Field	Description
Interface Name	Specify a numeric suffix for the interface (1-4999).
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Add an optional description of the interface.
Config Tab	
Virtual Router	Select a virtual router for this interface, or click New to configure a new virtual router. Refer to "Virtual Routers and Routing Protocols" on page 153. None removes the configuration from the interface.
Virtual System	Select the virtual system for the interface. None removes the configuration from the interface.

Table 59. Tunnel Interface Settings

Field	Description
Security Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151). None removes the configuration from the interface.
IPv4 Tab	
IP Address	Click Add to enter IP addresses and network masks for the interface.
IP∨6 Tab	
Enable IPv6 on the	Select the check box to enable IPv6 addressing for the interface.
interface	This option allows you to route IPv6 traffic over an IPv4 IPSec tunnel and will provide confidentiality between IPv6 networks. The IPv6 traffic is encapsulated by IPv4 and then ESP.
	To route IPv6 traffic to the tunnel, you will either use a static route to the tunnel, or use a Policy Based Forwarding (PBF) rule to direct traffic and to provide redundancy by monitoring the other end of the tunnel and failing over when needed.
Interface ID	Enter the 64-bit extended unique identifier in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. If the interface ID is left blank, the firewall will use the EUI-64 generated from the physical interface's MAC address.
Address	Click Add and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. Select Use interface ID as host portion to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address. Select Anycast to include routing through the nearest node. If Prefix is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.
Advanced Tab	
Other Info	Specify the following:
	• Management Profile—Select a profile that specifies which protocols, if any, can be used to manage the firewall over this interface.
	• MTU—Enter the MTU in bytes for packets sent on this interface (512- 1500, default 1500). If machines on either side of the firewall perform PMTUD, the MTU value will be returned in an ICMP fragmentation needed message indicating that the MTU is too large.
	Note: The firewall automatically considers tunnel overhead when performing IP fragmentation and also adjusts the TCP maximum segment size (MSS) as needed.

Table 59. Tunnel Interface Settings (Continued)

Configuring Tap Interfaces

► Network > Interfaces

You can define tap interfaces as needed to permit connection to a span port on a switch for traffic monitoring only (refer to "Tap Mode Deployments" on page 125).

To define tap interfaces, click an interface name on the **Ethernet** tab, and specify the following information.

Field	Description
Interface Name	Specify a name for the interface or keep the default name.
Interface Type	Select Tap from the drop-down list.
Netflow Profile	Select a profile if you want to export all ingress traffic through the interface to a specified NetFlow server. Refer to "Configuring Netflow Settings" on page 85.
Comment	Enter an optional description of the interface.
Config Tab	
Virtual System	Select a virtual system. None removes the configuration from the interface.
Zone	Select a security zone for the interface, or click New to define a new zone (refer to "Defining Security Zones" on page 151). None removes the configuration from the interface.
Advanced Tab	
Link Speed	Select the interface speed in Mbps (10, 100, or 1000) or auto.
Link Duplex	Select whether the interface transmission mode is full-duplex (full), half- duplex (half), or negotiated automatically (auto).
Link State	Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto).

Table 60. Tap Interface Settings

Configuring HA Interfaces

Each HA interface has a specific function: one interface is for configuration synchronization and heartbeats and the other interface is for state synchronization. If active/active high availability is enabled, a third HA interface can be used to forward packets.



Note: Some Palo Alto Networks firewalls include dedicated physical ports for use in HA deployments (one for the control link and one for the data link). For firewalls that do not include dedicated ports, you must specify the data ports that will be used for HA. For additional information on HA, refer to "Enabling HA on the Firewall" on page 101.

To define HA interfaces, click an interface name and specify the following information.

Field	Description
Interface Name	Choose the interface from the drop-down list. Modify the name if desired.
Interface Type	Select HA from the drop-down list.
Comment	Enter an optional description of the interface.
Advanced Tab	
Link Speed	Select the interface speed in Mbps (10, 100, or 1000) or auto.
Link Duplex	Select whether the interface transmission mode is full-duplex (full), half- duplex (half), or negotiated automatically (auto).
Link State	Select whether the interface status is enabled (up), disabled (down), or determined automatically (auto).

Table 61. HA Interface Settings

Security Zones

A security zone identifies one or more source or destination interfaces on the firewall. When you define a security policy rule, you must specify the source and destination security zones of the traffic. For example, an interface connected to the Internet is in an "untrusted" security zone, while an interface connected to the internal network is in a "trusted" security zone.

Separate zones must be created for each type of interface (Layer 2, Layer 3, tap, or virtual wire), and each interface must be assigned to a zone before it can process traffic. Security policies can be defined only between zones of the same type. However, if you create a VLAN interface for one or more VLANs, applying security policies between the VLAN interface zone and a Layer 3 interface zone (Figure 16) has the same effect as applying policies between the Layer 2 and Layer 3 interface zones.



Figure 16. Zone and Interface Types

Defining Security Zones

Network > Zones

In order for a firewall interface to be able to process traffic, it must be assigned to a security zone. To define security zones, click **New** and specify the following information:

Field	Description
Name	Enter a zone name (up to 15 characters). This name appears in the list of zones when defining security policies and configuring interfaces. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, periods, and underscores.
Location	Select the virtual system that applies to this zone.

Table 62. Security Zone Settings

Field	Description
Туре	Select a zone type (Layer2, Layer3, Virtual Wire, Tap, or External virtual system) to list all the interfaces of that type that have not been assigned to a zone. The Layer 2 and Layer 3 zone types list all Ethernet interfaces and subinterfaces of that type. The External virtual system type is for communications among virtual systems in the firewall. Refer to "Communications Among Virtual Systems" on page 111. Each interface can belong to one zone in one virtual system.
Zone Protection Profiles	Select a profile that specifies how the security gateway responds to attacks from this zone. To add new profiles, refer to "Defining Zone Protection Profiles" on page 178.
Log Setting	Select a log forwarding profile for forwarding zone protection logs to an external system.
Enable User Identification	Select to enable the user identification function on a per-zone basis.
User Identification ACL Include List	Enter the IP address or IP address/mask of a user or group to be identified (format <i>ip_address/mask</i> ; for example, 10.1.1.1/24). Click Add . Repeat as needed. If an include list is not configured, then all IP addresses are allowed.
User Identification ACL Exclude List	Enter the IP address or IP address/mask of a user or group that will explicitly not be identified (format <i>ip_address/mask</i> ; for example, 10.1.1.1/24). Click Add . Repeat as needed. If an exclude list is not configured, then all IP addresses are allowed.

Table 62. Security Zone Settings (Continued)

VLAN Support

► Network > VLANs

The firewall supports VLANs that conform to the IEEE 802.1Q standard. Each Layer 2 interface that is defined on the firewall must be associated with a VLAN. The same VLAN can be assigned to multiple Layer 2 interfaces, but each interface can belong to only one VLAN. Optionally, a VLAN can also specify a VLAN interface that can route traffic to Layer 3 destinations outside the VLAN.

Field	Description
Name	Enter a VLAN name (up to 31 characters). This name appears in the list of VLANs when configuring interfaces. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
VLAN Interface	Select a VLAN interface to allow traffic to be routed outside the VLAN. To define a VLAN interface, refer to "Configuring VLAN Interfaces" on page 143.
L3 Forwarding Enabled	If you select a VLAN interface, you can select the check box to enable Layer 3 routing over the selected interface.

Table 63. VLAN Settings

Field	Description
Interfaces	Specify firewall interfaces for the VLAN.
Static MAC Configuration	Specify the interface through which a MAC address is reachable. This will override any learned interface-to-MAC mappings.

Table 63.	VLAN Setti	ngs (Continued)
-----------	------------	-----------------

Virtual Routers and Routing Protocols

You can set up virtual routers to enable the firewall to route packets at Layer 3 by making packet forwarding decisions according to the destination IP address. The Ethernet interfaces and VLAN interfaces defined on the firewall receive and forward the Layer 3 traffic. The destination zone is derived from the outgoing interface based on the forwarding criteria, and policy rules are consulted to identify the security policies to be applied. In addition to routing to other network devices, virtual routers can route to other virtual routers within the same firewall if a next hop is specified to point to another virtual router.

Support is provided for static routing and dynamic routing using the Routing Information Protocol (RIP), Open Shortest Path First (OSPF) protocol, and Border Gateway Protocol (BGP).



Note: Policy-based forwarding is also supported for traffic on Layer 3 interfaces.

Routing Information Protocol

RIP was designed for small IP networks and relies on hop count to determine routes; the best routes have the fewest number of hops. RIP is based on UDP and uses port 520 for route updates. By limiting routes to a maximum of 15 hops, the protocol helps prevent the development of routing loops, but also limits the supported network size. If more than 15 hops are required, traffic is not routed. RIP also can take longer to converge than OSPF and other routing protocols. The firewall supports RIP v2.

Open Shortest Path First

OSPF determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The router keeps information about the links between it and the destination and can make highly efficient routing decisions. A cost is assigned to each router interface, and the best routes are determined to be those with the lowest costs, when summed over all the encountered outbound router interfaces and the interface receiving the LSA.

Hierarchical techniques are used to limit the number of routes that must be advertised and the associated LSAs. Because OSPF dynamically processes a considerable amount of route information, it has greater processor and memory requirements than does RIP.

Border Gateway Protocol

The Border Gateway Protocol (BGP) is the primary Internet routing protocol. BGP determines network reachability based on IP prefixes that are available within autonomous systems (AS), where an AS is a set of IP prefixes that a network provider has designated to be part of a single routing policy.

In the routing process, connections are established between BGP peers (or neighbors). If a route is permitted by the policy, it is stored in the routing information base (RIB). Each time the local firewall RIB is updated, the firewall determines the optimal routes and sends an update to the external RIB, if export is enabled.

Conditional advertisement is used to control how BGP routes are advertised. The BGP routes must satisfy conditional advertisement rules before being advertised to peers.

BGP supports the specification of aggregates, which combine multiple routes into a single route. During the aggregation process, the first step is to find the corresponding aggregation rule by performing a longest match that compares the incoming route with the prefix values for other aggregation rules.

The firewall provides a complete BGP implementation that includes the following features:

- Specification of one BGP routing instance per virtual router.
- Routing policies based on route-map to control import, export and advertisement, prefixbased filtering, and address aggregation.
- Advanced BGP features that include route reflector, AS confederation, route flap dampening, and graceful restart.
- IGP-BGP interaction to inject routes to BGP using redistribution profiles.

BGP configuration consists of the following elements:

- Per-routing-instance settings, which include basic parameters such as local route ID and local AS and advanced options such as path selection, route reflector, AS confederation, route flap, and dampening profiles.
- Authentication profiles, which specify the MD5 authentication key for BGP connections.
- Peer group and neighbor settings, which include neighbor address and remote AS and advanced options such as neighbor attributes and connections.
- Routing policy, which specifies rule sets that peer groups and peers use to implement imports, exports, conditional advertisements, and address aggregation controls.

Multicast Routing

The multicast routing feature allows the firewall to route multicast streams using Protocol Independent Multicast Sparse Mode (PIM-SM) and PIM Source Specific Multicast (PIM-SSM) for applications such as media broadcasting (radio and video) with PIMv2. The firewall performs Internet Group Management Protocol (IGMP) queries for hosts that are on the same network as the interface on which IGMP is configured. PIM-SM and IGMP can be enabled on Layer 3 interfaces. IGMP v1, v2, and v3 are supported. PIM and IGMP must be enabled on host-facing interfaces.

PAN-OS provides full multicast security while acting as a PIM designated router (DR), PIM rendezvous point (RP), intermediate PIM router, or IGMP querier. The firewall can be deployed in environments in which the RP is statically configured or dynamically elected. The

bootstrap router (BSR) role is not supported. Deployment across IPSec tunnels is fully supported between Palo Alto Networks firewalls. GRE encapsulation within IPSec is not currently supported.

Security policy

PAN-OS provides two methods to enforce security on multicast feeds. Multicast groups can be filtered in the IGMP and PIM group permission settings specified on an interface level. Multicast traffic must also be explicitly allowed by security policy. A special destination zone known as "Multicast" has been added and must be specified to control multicast traffic in security, QoS, and DoS protection rules. In contrast to unicast security policy, multicast security policies must be explicitly created when the source and destination interfaces are in the same zone. Security profiles are supported in multicast environments that require threat prevention capabilities.

Logging

Each multicast session passing through the firewall creates only one traffic log entry (even if the firewall is replicating packets for distribution on multiple interfaces). Traffic logs indicate the number of bytes coming into the firewall rather than the number of bytes distributed as part of the multicast feed.

Defining Virtual Routers

Network > Virtual Routers

Defining virtual routers allows you to set up forwarding rules for Layer 3 and enable the use of dynamic routing protocols. Each Layer 3 interface, loopback interface, and VLAN interface defined on the firewall should be associated with a virtual router. Each interface can belong to only one virtual router.



Note: To configure Ethernet ports as Layer 3 interfaces, refer to "Configuring Layer 3 Interfaces" on page 130. To define Layer 3 subinterfaces, refer to "Configuring Layer 3 Subinterfaces" on page 134. For an overview of virtual routers, refer to "Virtual Routers and Routing Protocols" on page 153.

Define settings on the specified tabs, as appropriate.

General Tab

Select the interfaces to include in the virtual router and add any static routes. Refer to the following table.

Field	Description
Name	Specify a name to describe the virtual router (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Interfaces	Select the interfaces that you want to include in the virtual router. When you select an interface, it is included in the virtual router and can be used as an outgoing interface in the virtual router's routing tab.
	To specify the interface type, refer to "Firewall Interfaces" on page 127.
	<i>Note:</i> When you add an interface, its connected routes are added automatically.

Table 64. Virtual Router Settings - General Tab

Field	Description
Administrative	Specify the following administrative distances:
Distances	• Static routes (10-240, default 10).
	• OSPF Int (10-240, default 30).
	• OSPF Ext (10-240, default 110).
	• IBGP (10-240, default 200).
	• EBGP (10-240, default 20).
	• RIP (10-240, default 120).

Table 64. Virtual Router Settings - General Tab (Continued)

Static Routes Tab

Optionally enter one or more static routes. Click the IP or IPv6 tab to specify the route using IPv4 or IPv6 addresses. It is usually necessary to configure default routes (0.0.0.0/0) here. Default routes are applied for destinations that are otherwise not found in the virtual router's routing table.

Field	Description
Name	Enter a name to identify the static route (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Destination	Enter an IP address and network mask in the format <i>ip_address/mask</i> .
Interface	Select the interface to forward packets to the destination, or configure the next hop settings, or both.
Next Hop	Specify the following next hop settings:
	• None—Select if there is no next hop for the route.
	• IP Address—Specify the IP address of the next hop router.
	• Discard —Select if you want to drop traffic that is addressed to this des- tination.
	• Next VR—Select a virtual router in the firewall as the next hop. This option allows you to route internally between virtual routers within a single firewall.
Admin Distance	Specify the administrative distance for the static route (10-240, default 10).
Metric	Specify a valid metric for the static route (1 - 65535).
No Install	Select if you do not want to install the route in the forwarding table. The route is retained in the configuration for future reference.

Table 65. Virtual Router Settings - Static Routes Tab

Redistribution Profiles Tab

Modify route redistribution filter, priority, and action based on desired network behavior. Route redistribution allows static routes and routes that are acquired by other protocols to be advertised through specified routing protocols. Redistribution profiles must be applied to routing protocols in order to take effect. Without redistribution rules, each protocol runs separately and does not communicate outside its purview. Redistribution profiles can be added or modified after all routing protocols are configured and the resulting network topology is established. Apply redistribution profiles to the RIP and OSPF protocols by defining export rules. Apply redistribution profiles to BGP in the Redistribution Rules tab. Refer to the following table.

Field	Description
Name	Click Add to display the Redistribution Profile page, and enter the profile name.
Priority	Enter a priority (range 1-255) for this profile. Profiles are matched in order (lowest number first).
Redistribute	Choose whether to perform route redistribution based on the settings in this window.
	• Redist —Select to redistribute matching candidate routes. If you select this option, enter a new metric value. A lower metric value means a more preferred route.
	• No Redist—Select to not redistribute matching candidate routes.
General Filter Tab	
Туре	Select check boxes to specify the route types of the candidate route.
Interface	Select the interfaces to specify the forwarding interfaces of the candidate route.
Destination	To specify the destination of the candidate route, enter the destination IP address or subnet (format $x.x.x.x$ or $x.x.x/n$) and click Add . To remove an entry, click the \Box icon associated with the entry.
Next Hop	To specify the gateway of the candidate route, enter the IP address or subnet (format x.x.x.x or x.x.x/n) that represents the next hop and click Add . To remove an entry, click the \bigcirc icon associated with the entry.
OSPF Filter Tab	
Path Type	Select check boxes to specify the route types of the candidate OSPF route.
Area	Specify the area identifier for the candidate OSPF route. Enter the OSPF area ID (format x.x.x.x), and click Add . To remove an entry, click the c icon associated with the entry.
Tag	Specify OSPF tag values. Enter a numeric tag value (1-255), and click Add. To remove an entry, click the \Box icon associated with the entry.
BGP Filter Tab	
Community	Specify a community for BGP routing policy.
Extended Community	Specify an extended community for BGP routing policy.

Table 66. Virtual Router Settings - Redistribution Profiles Tab

RIP Tab

Specify parameters for use of the Routing Information Protocol (RIP) on the selected interfaces. Although it is possible to configure both RIP and OSPF, it is generally recommended to choose only one of these protocols. Refer to the following table.

Field	Description
Enable	Select the check box to enable the RIP protocol.
Reject Default Route	Select the check box if you do not want to learn any default routes through RIP. Selecting the check box is highly recommended.
Allow Redist Default Route	Select the check box to permit redistribution of default routes through RIP.
Interfaces	
Interface	Select the interface that runs the RIP protocol.
Enable	Select to enable these settings.
Advertise	Select to advertise a default route to RIP peers with the specified metric value.
Metric	Specify a metric value for the router advertisement. This field is visible only if the Advertise check box is selected.
Auth Profile	Select the profile.
Mode	Select normal , passive , or send-only .
Timers	
Interval Seconds (sec)	Define the length of the timer interval in seconds. This duration is used for the remaining RIP timing fields (1 - 60).
Update Intervals	Enter the number of intervals between route update announcements (1 - 3600).
Expire Intervals	Enter the number of intervals between the time that the route was last updated to its expiration (1- 3600).
Delete Intervals	Enter the number of intervals between the time that the route expires to its deletion (1- 3600).
Auth Profiles	
Profile Name	Enter a name for the authentication profile to authenticate RIP messages. To authenticate RIP messages, first define the authentication profiles and then apply them to interfaces on the RIP tab.
Password Type	Select the type of password (simple or MD5).
	• If you select Simple , enter the simple password and then confirm.
	• If you select MD5 , enter one or more password entries, including Key- ID (0-255), Key , and optional Preferred status. Click Add for each entry, and then click OK . To specify the key to be used to authenticate outgoing message, select the Preferred option.

Table 67. Virtual Router Settings - RIP Tab

Field	Description
Export Rules	
Export Rules	(Read-only) Displays the rules that apply to routes sent by the virtual router to a receiving router.
	• Allow Redistribute Default Route—Select the check box to permit the firewall to redistribute its default route to peers.
	• Redistribution Profile —Select a redistribution profile that allows you to modify route redistribution, filter, priority, and action based on the desired network behavior. Refer to "Redistribution Profiles Tab" on page 158.

Table 67. Virtual Router Settings - RIP Tab (Continued)

OSPF Tab

Specify parameters for use of the Open Shortest Path First (OSPFv2) protocol on the selected interfaces. Although it is possible to configure both RIP and OSPF, it is generally recommended to choose only one IGP protocol. Refer to the following table.BGP Tab

Field	Description
Enable	Select the check box to enable the OSPF protocol.
Reject Default Route	Select the check box if you do not want to learn any default routes through OSPF. Selecting the check box is recommended, especially for static routes.
Router ID	Specify the router ID associated with the OSPF instance in this virtual router. The OSPF protocol uses the router ID to uniquely identify the OSPF instance.
Areas	
Area ID	Configure the area over which the OSPF parameters can be applied.
	Enter an identifier for the area in x.x.x.x format. This is the identifier that each neighbor must accept to be part of the same area.

Table 68. Virtual Router Settings - OSPF Tab

Field	Description
Туре	Select one of the following options.
	• Normal—There are no restrictions; the area can carry all types of routes.
	 Stub—There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, select Accept Summary if you want to accept this type of link state advertisement (LSA) from other areas. Also, specify whether to include a default route LSA in advertisements to the stub area along with the associated metric value (1-255). If the Accept Summary option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs. NSSA (Not-So-Stubby Area)—It is possible to leave the area directly, but only by routes other than OSPF routes. If you select this option, select Accept Summary if you want to accept this type of LSA. Specify whether to include a default route LSA in advertisements to the stub
	area along with the associated metric value (1-255). Also, select the route type used to advertise the default LSA. Click Add in the External Ranges section and enter ranges if you want to enable or suppress advertising external routes that are learned through NSSA to other areas.
Range	Click Add to aggregate LSA destination addresses in the area into subnets. Enable or suppress advertising LSAs that match the subnet, and click OK . Repeat to add additional ranges.
Interface	Click Add and enter the following information for each interface to be included in the area, and click OK .
	• Interface—Choose the interface.
	• Enable—Cause the OSPF interface settings to take effect.
	• Passive —Select the check box to if you do not want the OSPF interface to send or receive OSPF packets. Although OSPF packets are not sent or received if you choose this option, the interface is included in the LSA database.
	• Link type—Choose Broadcast if you want all neighbors that are accessible through the interface to be discovered automatically by multicasting OSPF hello messages, such as an Ethernet interface. Choose p2p (point-to-point) to automatically discover the neighbor. Choose p2mp (point-to-multipoint) when neighbors must be defined manually. Defining neighbors manually is allowed only for p2mp mode.
	• Metric—Enter the OSPF metric for this interface (0-65535).
	• Priority —Enter the OSPF priority for this interface (0-255). It is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR) according to the OSPF protocol. When the value is zero, the router will not be elected as a DR or BDR.
	Auth Profile—Select a previously-defined authentication profile.
	 Timing—It is recommended that you keep the default timing settings. Neighbors—For p2pmp interfaces, enter the neighbor IP address for all neighbors that are reachable through this interface.

Table 68. Virtual Router Settings - OSPF Tab (Continued)

Field	Description
Virtual Link	Configure the virtual link settings to maintain or enhance backbone area connectivity. The settings must be defined for area boarder routers, and must be defined within the backbone area (0.0.0.0). Click Add , enter the following information for each virtual link to be included in the backbone area, and click OK .
	• Name—Enter a name for the virtual link.
	• Neighbor ID —Enter the router ID of the router (neighbor) on the other side of the virtual link.
	• Transit Area —Enter the area ID of the transit area that physically contains the virtual link.
	• Enable—Select to enable the virtual link.
	• Timing—It is recommended that you keep the default timing settings.
	• Auth Profile—Select a previously-defined authentication profile.
Auth Profiles	
Profile Name	Enter a name for the authentication profile. To authenticate the OSPF messages, first define the authentication profiles and then apply them to interfaces on the OSPF tab.
Password Type	Select the type of password (simple or MD5).
	• If you select Simple , enter the password.
	• If you select MD5 , enter one or more password entries, including Key- ID (0-255), Key , and optional Preferred status. Click Add for each entry, and then click OK . To specify the key to be used to authenticate outgoing message, select the Preferred option.
Export Rules	
Allow Redistribute Default Route	Select the check box to permit redistribution of default routes through OSPF.
Name	Select the name of a redistribution profile. The value must be an IP subnet or valid redistribution profile name.
New Path Type	Choose the metric type to apply.
New Tag	Specify a tag for the matched route that has a 32-bit value.
Advanced	
RFC 1583 Compatibility	Select the check box to assure compatibility with RFC 1583.
Timers	 SPF Calculation Delay (sec)—This option is a delay timer allowing you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times. LSA Interval (sec)—The option specifies the minimum time between
	transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur.

Table 68. Virtual Router Settings - OSPF Tab (Continued)

BGP Tab

Specify parameters for use of Border Gateway Protocol (BGP) on the selected interfaces. Refer to the following table.

	•
Field	Description
Enable	Select the check box to enable BGP.
Router ID	Enter the IP address to assign to the virtual router.
AS Number	Enter the number of the AS to which the virtual router belongs, based on the router ID (range 1-4294967295).
General Tab	
Reject Default Route	Select the check box to ignore any default routes that are advertised by BGP peers.
Install Route	Select the check box to install BGP routes in the global routing table.
Aggregate MED	Select to enable route aggregation even when routes have different Multi- Exit Discriminator (MED) values.
Default Local Preference	Specifies a value than can be used to determine preferences among different paths.
AS Format	Select the 2-byte (default) or 4-byte format. This setting is configurable for interoperability purposes.
Always Compare MED	Enable MED comparison for paths from neighbors in different autonomous systems.
Deterministic MED Comparison	Enable MED comparison to choose between routes that are advertised by IBGP peers (BGP peers in the same autonomous system).
Auth Profiles	Click Add to include a new authentication profile and configure the following settings:
	• Profile Name —Enter a name to identify the profile.
	 Secret/Confirm Secret—Enter and confirm a passphrase for BGP peer communications.
	Click the 🖻 icon to delete a profile.
Advanced Tab	
Graceful Restart	Activate the graceful restart option.
	• Stale Route Time—Specify the length of time that a route can stay in the stale state (range 1-3600 seconds, default 120 seconds).
	• Local Restart Time—Specify the length of time that the local device takes to restart. This value is advertised to peers (range 1-3600 seconds, default 120 seconds).
	• Max Peer Restart Time—Specify the maximum length of time that the local device accepts as a grace period restart time for peer devices (range 1-3600 seconds, default 120 seconds).
Reflector Cluster ID	Specify an IPv4 identifier to represent the reflector cluster.
Confederation Member AS	Specify the identifier for the AS confederation to be presented as a single AS to external BGP peers.

Table 69. Virtual Router Settings - BGP Tab

Field	Description
Dampening Profiles	Settings include:
	• Profile Name —Enter a name to identify the profile.
	• Enable—Activate the profile.
	 Cutoff—Specify a route withdrawal threshold above which a route advertisement is suppressed (range 0.0-1000.0, default 1.25).
	 Reuse—Specify a route withdrawal threshold below which a suppressed route is used again (range 0.0-1000.0, default 5).
	• Max. Hold Time—Specify the maximum length of time that a route can be suppressed, regardless of how unstable it has been (range 0-3600 seconds, default 900 seconds).
	• Decay Half Life Reachable—Specify the length of time after which a route's stability metric is halved if the route is considered reachable (range 0-3600 seconds, default 300 seconds).
	• Decay Half Life Unreachable —Specify the length of time after which a route's stability metric is halved if the route is considered unreachable (range 0-3600 seconds, default 300 seconds).
	Click the 🖻 icon to delete a profile.
Peer Group Tab	
Name	Enter a name to identify the peer.
E 11	

 Table 69.
 Virtual Router Settings - BGP Tab (Continued)

Name	Enter a name to identify the peer.			
Enable	Select to activate the peer.			
Aggregated Confed AS Path	Select the check box to include a path to the configured aggregated confederation AS.			
Soft Reset with Stored Info	Select the check box to perform a soft reset of the firewall after updating the peer settings.			
Туре	Specify the type of peer or group and configure the associated settings (see below in this table for descriptions of Import Next Hop and Export Next Hop).			
	• IBGP —Specify the following;			
	– Export Next Hop			
	• EBGP Confed—Specify the following;			
	– Export Next Hop			
	• IBGP Confed —Specify the following;			
	– Export Next Hop			
	• EBGP —Specify the following:			
	– Import Next Hop			
	– Export Next Hop			
	 Remove Private AS (select if you want to force BGP to remove private AS numbers). 			
Import Next Hop	Choose an option for next hop import:			
	 original—Use the Next Hop address provided in the original route advertisement. 			
	• use-peer —Use the peer's IP address as the Next Hop address.			

Field	Description		
Export Next Hop	Choose an option for next hop export:		
	• resolve —Resolve the Next Hop address using the local forwarding table.		
	• use-self —Replace the Next Hop address with this router's IP address to ensure that it will be in the forwarding path.		
Peer	To add a new peer, click New and configure the following settings:		
	• Name—Enter a name to identify the peer.		
	• Enable—Select to activate the peer.		
	• Peer AS —Specify the AS of the peer.		
	• Local Address—Choose a firewall interface and local IP address.		
	• Connection Options—Specify the following options:		
	- Auth Profile—Select the profile.		
	 Keep Alive Interval—Specify an interval after which routes from a peer are suppressed according to the hold time setting (range 0-1200 seconds, default 30 seconds). 		
	 Multi Hop—Set the time-to-live (TTL) value in the IP header (range 1-255, default 0). The default value of 0 means 2 for eBGP and 255 for iBGP. 		
	 Open Delay Time—Specify the delay time between opening the peer TCP connection and sending the first BGP open message (range 0-240 seconds, default 0 seconds). 		
	 Hold Time—Specify the period of time that may elapse between successive KEEPALIVE or UPDATE messages from a peer before the peer connection is closed. (range 3-3600 seconds, default 90 seconds). 		
	 - Idle Hold Time—Specify the time to wait in the idle state before retrying connection to the peer (range 1-3600 seconds, default 15 seconds). 		
	• Peer Address —Specify the IP address and port of the peer.		
	• Advanced Options—Configure the following settings:		
	 Reflector Client—Select the type of reflector client (Non-Client, Client, or Meshed Client). Routes that are received from reflector clients are shared with all internal and external BGP peers. 		
	- Peering Type —Specify a bilateral peer, or leave unspecified.		
	 Max. Prefixes—Specify the maximum number of supported IP prefixes (1 - 100000 or unlimited). 		
	• Incoming Connections/Outgoing Connections —Specify the incoming and outgoing port numbers and select the Allow check box to allow traffic to or from these ports.		

Table 69. Virtual Router Settings - BGP Tab (Continued)

Field	Description		
Import Rules/Export Rules Tabs			
Import Rules/Export Rules	Click the BGP Import Rules or Export Rules subtab. To add a new rule, click Add and configure the following settings.		
	• General subtab:		
	 Name—Specify a name to identify the rule. 		
	– Enable —Select to activate the rule.		
	 Used by—Select the peer groups that will use this rule. 		
	• Match subtab:		
	 AS-Path Regular Expression—Specify a regular expression for filtering of AS paths. 		
	 Community Regular Expression—Specify a regular expression for filtering of community strings. 		
	 Extended Community Regular Expression—Specify a regular expression for filtering of extended community strings. 		
	- Address Prefix—Specify IP addresses or prefixes for route filtering.		
	 MED—Specify a MED value for route filtering. 		
	- Next Hop—Specify next hop routers or subnets for route filtering.		
	 From Peer—Specify peer routers for route filtering. 		
	• Action subtab:		
	 Action—Specify an action (Allow or Deny) to take when the match conditions are met. 		
	 Local Preference—Specify a local preference metric, only if the action is Allow. 		
	- MED—Specify a MED value, only if the action is Allow (0- 65535).		
	 Weight—Specify a weight value, only if the action is Allow (0- 65535). 		
	- Next Hop—Specify a next hop router, only if the action is Allow.		
	 Origin—Specify the path type of the originating route: IGP, EGP, or incomplete, only if the action is Allow. 		
	- AS Path Limit —Specify an AS path limit, only if the action is Allow .		
	 AS Path—Specify an AS path: None, Remove, Prepend, Remove and Prepend, only if the action is Allow. 		
	 Community—Specify a community option: None, Remove All, Remove Regex, Append, or Overwrite, only if the action is Allow. 		
	 Extended Community—Specify a community option: None, Remove All, Remove Regex, Append, or Overwrite, only if the action is Allow. 		
	 Dampening—Specify the dampening parameter, only if the action is Allow. 		
	Click the 🖃 icon to delete a group. Click Clone to add a new group with the same settings as the selected group. A suffix is added to the new group name to distinguish it from the original group.		

Table 69. Virtual Router Settings - BGP Tab (Continued)

Field	Description			
Conditional Adv Tab	The BGP conditional advertisement feature allows you to control what route to advertise in the event that a different route is not available in the local BGP routing table (LocRIB), indicating a peering or reachability failure. This is useful in cases where you want to try and force routes to one AS over another, for example if you have links to the Internet through multiple ISPs and you want traffic to be routed to one provider instead of the other unless there is a loss of connectivity to the preferred provider. With conditional advertising, you can configure a non-exist filter that matches the prefix of the preferred route. If any route matching the non- exist filter is not found in the local BGP routing table, only then will the device allow advertisement of the alternate route (the route to the other, non-preferred provider) as specified in its advertise filter. To configure conditional advertisement, select the Conditional Adv tab and then click Add . The following describes how to configure the values in the fields.			
Policy	Specify the policy name for this conditional advertisement rule.			
Enable	Select the check box to enable BGP conditional advertisement.			
Used By	Click Add and select the peer groups that will use this conditional advertisement policy.			
Non Exist Filters Subtab	Use this tab to specify the prefix(es) of the preferred route. This specifies the route that you want to advertise, if it is available in the local BGP routing table. If a prefix is going to be advertised and matches a Non Exist filter, the advertisement will be suppressed.			
	Click Add to create a non-exist filter.			
	• Non Exist Filters—Specify a name to identify this filter.			
	• Enable—Select to activate the filter.			
	• AS-Path Regular Expression—Specify a regular expression for filtering of AS paths.			
	• Community Regular Expression —Specify a regular expression for fil- tering of community strings.			
	• Extended Community Regular Expression—Specify a regular expression for filtering of extended community strings.			
	• MED—Specify a MED value for route filtering.			
	• Address Prefix—Click Add and then specify the exact NLRI prefix for the preferred route.			
	• Next Hop—Specify next hop routers or subnets for route filtering.			
	 From Peer—Specify peer routers for route filtering. 			

Table 69. Virtual Router Settings - BGP Tab (Continued)

Description	
Use this tab to specify the prefix(es) of the route in the Local-RIB routing table that should be advertised in the event that the route in the non-exist filter is not available in the local routing table.	
If a prefix is going to be advertised and does not match a Non Exist filter, the advertisement will occur.	
Click Add to create an advertise filter.	
• Advertise Filters—Specify a name to identify this filter.	
• Enable—Select to activate the filter.	
• AS-Path Regular Expression —Specify a regular expression for filtering of AS paths.	
• Community Regular Expression —Specify a regular expression for filtering of community strings.	
• Extended Community Regular Expression—Specify a regular expression for filtering of extended community strings.	
• MED—Specify a MED value for route filtering.	
• Address Prefix—Click Add and then specify the exact NLRI prefix for the route to be advertised if the preferred route is not available.	
• Next Hop—Specify next hop routers or subnets for route filtering.	
• From Peer—Specify peer routers for route filtering.	
To add a new rule, click Add to display the settings. Configure the settings in the Suppress Filters , Advertise Filters , and Aggregate Route Attributes subtabs, and click Done to add the rule to the Addresses list. The parameters are described above in this table for the Import Rules and Export Rules tabs.	
Click the 🖃 icon to delete a rule.	
Select the name of a redistribution profile. The value must be an IP subnet or valid redistribution profile name.	
Select the check box to permit the firewall to redistribute its default route to BGP peers.	
To add a new rule, click Add , configure the settings, and click Done . The parameters are described above in this table for the Import Rules and Export Rules tabs.	

 Table 69.
 Virtual Router Settings - BGP Tab (Continued)

Multicast Tab

Specify settings for multicast routing in the following table.

Field	Description			
Enable	Select the check box to enable multicast routing.			
Rendezvous Point Subtab				
RP Type	Choose the type of Rendezvous Point (RP) that will run on this virtual router. A static RP must be explicitly configured on other PIM routers whereas a candidate RP is elected automatically.			
	• None—Choose if there is no RP running on this virtual router.			
	• Static—Specify a static IP address for the RP and choose options for RP Interface and RP Address from the drop-down lists. Select the Over- ride learned RP for the same group check box if you want to use the specified RP instead of the RP elected for this group.			
	• Candidate —Specify the following information for the candidate RP running on this virtual router:			
	 RP Interface—Select an interface for the RP. Valid interface types include loopback, L3, VLAN, aggregate Ethernet, and tunnel. 			
	- RP Address —Select an IP address for the RP.			
	- Priority —Specify a priority for candidate RP messages (default 192).			
	 Advertisement interval—Specify an interval between advertise- ments for candidate RP messages. 			
	• Group list—If you choose Static or Candidate, click Add to specify a list of groups for which this candidate RP is proposing to be the RP.			
Remote Rendezvous	Click Add and specify the following:			
Point	• IP address —Specify the IP address for the RP.			
	• Override learned RP for the same group —Select the check box to use the specified RP instead of the RP elected for this group.			
	• Group —Specify a list of groups for which the specified address will act as the RP.			
Interfaces Subtab				
Name	Enter a name to identify an interface group.			
Description	Enter an optional description.			
Interface	Click Add to specify one or more firewall interfaces.			
Group Permissions	Specify general rules for multicast traffic:			
*	• Any Source —Click Add to specify a list of multicast groups for which PIM-SM traffic is permitted.			
	 Source-Specific—Click Add to specify a list of multicast group and multicast source pairs for which PIM-SSM traffic is permitted. 			

Table 70. Virtual Router Settings - Multicast Tab

Field	Description	
IGMP	Specify rules for IGMP traffic. IGMP must be enabled for host facing interfaces (IGMP router) or for IGMP proxy host interfaces.	
	• Enable—Select the check box to enable the IGMP configuration.	
	• IGMP Version—Choose version 1, 2, or 3 to run on the interface.	
	• Enforce Router-Alert IP Option—Select the check box to require the router-alert IP option when speaking IGMPv2 or IGMPv3. This option must be disabled for compatibility with IGMPv1.	
	• Robustness —Choose an integer value to account for packet loss on a network (range 1-7, default 2). If packet loss is common, choose a higher value.	
	• Max Sources —Specify the maximum number of source-specific memberships allowed on this interface (0 = unlimited).	
	• Max Groups —Specify the maximum number of groups allowed on this interface.	
	• Query Configuration—Specify the following:	
	 Query interval—Specify the interval at which general queries are sent to all hosts. 	
	 Max Query Response Time—Specify the maximum time between a general query and a response from a host. 	
	 Last Member Query Interval—Specify the interval between group or source-specific query messages (including those sent in response to leave-group messages). 	
	 Immediate Leave—Select the check box to leave the group immedi- ately when a leave message is received. 	
PIM configuration	Specify the following Protocol Independent Multicast (PIM) settings:	
	• Enable—Select the check box to allow this interface to receive and/or forward PIM messages	
	• Assert Interval—Specify the interval between PIM assert messages.	
	• Hello Interval—Specify the interval between PIM hello messages.	
	• Join Prune Interval—Specify the interval between PIM join and prune messages (seconds). Default is 60.	
	• DR Priority—Specify the designated router priority for this interface	
	• BSR Border —Select the check box to use the interface as the bootstrap border.	
	• PIM Neighbors —Click Add to specify the list of neighbors that will communicate with using PIM.	
SPT Threshold Subtab		
Name	The Shortest Path Tree (SPT) threshold defines the throughput rate (in kbps) at which multicast routing will switch from shared tree distribution (sourced from the rendezvous point) to source tree distribution.	
	Click Add to specify the following SPT settings:	
	• Multicast Group Prefix —Specify the multicast IP address/prefix for which the SPT will be switched to source tree distribution when the throughput reaches the desired threshold (kbps).	
	• Threshold —Specify the throughput at which we'll switch from shared tree distribution to source tree distribution	

Table 70. Virtual Router Settings - Multicast Tab (Continued)

Field	Description		
Source Specific Address Space Subtab			
Name	Defines the multicast groups for which the firewall will provide source- specific multicast (SSM) services.		
	Click Add to specify the following settings for source-specific addresses:		
	• Name—Enter a name to identify this group of settings.		
	 Group—Specify groups for the SSM address space. 		
	• Included —Select this check box to include the specified groups in the SSM address space.		

Table 70. Virtual Router Settings - Multicast Tab (Continued)

Displaying Runtime Statistics for Virtual Routers

Network > Virtual Routers

Detailed runtime statistics are available for the virtual router and dynamic routing protocols from the **Virtual Routers** page. In the **Runtime Stats** column, click the **More Runtime Stats** link to open a new window that contains the routing table as well as routing protocol-specific details. For an overview of virtual routers, refer to "Virtual Routers and Routing Protocols" on page 153.

DHCP Server and Relay

Network > DHCP

The firewall supports the selection of DHCP servers or DHCP relay for IP address assignment on the Layer 3 and VLAN interfaces. Multiple DHCP servers are supported. Client requests can be forwarded to all servers, with the first server response sent back to the client.

The DHCP assignment also works across an IPSec VPN, allowing clients to receive an IP address assignment from a DHCP server on the remote end of an IPSec tunnel. For information on IPSec VPN tunnels, refer to "Configuring IPSec Tunnels" on page 309.

The settings depend on whether you select DHCP Server or DHCP Relay as the type.

Field	Description		
DHCP Server Tab			
Interface	Select the firewall interface.		
Mode	Choose whether the settings on this page are enabled, disabled, or are determined automatically.		
Ping IP when allocating new IP	Select the check box to send a ping message when allocating a new IP address.		

 Table 71.
 DHCP Settings

Field	Description			
Lease	Select Unlimited , or select Timeout and enter any limitations on the DHCP lease interval. You can enter days, hours, or minutes. For example, if you enter only hours, then the lease is restricted to that number of hours.			
Inheritance Source	Select a source to propagate various server settings from a DHCP client interface or PPPoE client interface into the DHCP server. Once an inheritance source is specified, select the desired fields in the DHCP server configuration and set to inherited to inherit the values from the DHCP server. For example, you can inherit DNS, WINS, NIS, NTP.			
Primary DNS Secondary DNS	Enter the IP address of the preferred and alternate Domain Name System (DNS) servers. The alternate server address is optional.			
Primary WINS Secondary WINS	Enter the IP address of the preferred and alternate Windows Internet Naming Service (WINS) servers. The alternate server address is optional.			
Primary NIS Secondary NIS	Enter the IP address of the preferred and alternate Network Information Service (NIS) servers. The alternate server address is optional.			
Primary NTP Secondary NTP	Enter the IP address of the preferred and alternate Network Time Protocol server. The alternate server address is optional.			
Gateway	Enter the IP address of the network gateway that is used to reach the DHCP servers.			
POP3 Server	Enter the IP address of the Post Office Protocol (POP3) server.			
SMTP Server	Enter the IP address of the Simple Mail Transfer Protocol (SMTP) server.			
DNS Suffix	Enter a suffix for the client to use locally when an unqualified hostname is entered that it cannot resolve.			
IP Pools	Specify the range of IP addresses to which this DHCP configuration applies and click Add . You can enter an IP subnet and subnet mask (for example, 192.168.1.0/24) or a range of IP addresses (for example, 192.168.1.10-192.168.1.20). Add multiple entries to specify multiple IP address pools.			
	To edit an existing entry, click Edit , make the changes, and click Done . To delete an entry, click Delete .			
	<i>Note: If you leave this area blank, there will be no restrictions on the IP ranges.</i>			
Reserved Address	Specify the IP address (format <i>x.x.x.x</i>) or MAC address (format <i>xx:xx:xx:xx:xx</i>) of any devices that you do not want to subject to DHCP address assignment.			
	To edit an existing entry, click Edit , make the changes, and click Done . To delete an entry, click Delete .			
	<i>Note:</i> If you leave this area blank, then there will be no reserved IP addresses.			
DHCP Relay Tab				
Interface	Select the firewall interface.			
IPv4	Select the IPv4 check box to use IPv4 addresses for DHCP relay and specify IPv4 addresses for up to four DHCP servers.			
IPv6	Select the IPv6 check box to use IPv6 addresses for DHCP relay and specify IPv6 addresses for up to four DHCP servers. Specify an outgoing interface if you are using an IPv6 multicast address for your server.			

Table 71. DHCP Settings (Continued)

DNS Proxy

Network > DNS Proxy

For all DNS queries that are directed to an interface IP address, the firewall supports the selective directing of queries to different DNS servers based on full or partial domain names. TCP or UDP DNS queries are sent through the configured interface. UDP queries switch over to TCP when a DNS query answer is too long for a single UDP packet.

If the domain name is not found in the DNS proxy cache, the domain name is searched for a match based on configuration of the entries in the specific DNS proxy object (on the interface on which the DNS query arrived) and forwarded to a name server based on the match results. If no match is found, the default name servers are used. Static entries and caching are also supported.

Field	Description			
Name	Specify a name to identify the DNS proxy (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.			
Enable	Select the check box to enable DNS proxy.			
Inheritance Source	Select a source to inherit default DNS server settings. This is commonly used in branch office deployments where the firewall's WAN interface is addressed by DHCP or PPPoE.			
Primary Secondary	Specify the IP addresses of the default primary and secondary DNS servers. If the primary DNS server cannot be found, the secondary will be used.			
Check inheritance source	Click the link to see the server settings that are currently assigned to the DHCP client and PPPoE client interfaces. These may include DNS, WINS, NTP, POP3, SMTP, or DNS suffix.			
Interface	Select the Interface check box to specify the firewall interfaces to support the DNS proxy rules. Select an interface from the drop-down list and click Add . You can add multiple interfaces. To delete an interface, select the interface and click Delete .			
DNS Proxy Rules	Identify DNS proxy server rules. Click Add and specify the following information:			
	• Name —A name is required so that a static entry can be reference and modified via the CLI.			
	• Turn on caching of domains resolved by this mapping —Select the check box to enable caching of domains that are resolved by this mapping.			
	• Domain Name—Click Add and enter the proxy server domain name. Repeat to add additional names. To delete a name, select the name and click Delete. For a DNS proxy rule, the number of tokens in a wildcard string must match the number of tokens in the requested domain. For example, "*.engineering.local" will not match "engineering.local". Both must be specified.			
	• Primary/Secondary —Enter the hostname or IP addresses of the primary and secondary DNS servers.			

Table 72.	DNS F	roxy	Settings
-----------	-------	------	----------

Field	Description
Static Entries	Provide static FQDN to IP address mappings that will be delivered in response to DNS queries made by hosts. Click Add and specify the following information:
	• Name—Enter a name for the Static Entry.
	 FQDN—Enter the Fully Qualified Domain Name (FQDN) that will be mapped to the static IP addresses defined in the Address field.
	 Address—Click Add and enter the IP addresses that map to this domain.
	Repeat to add additional addresses. To delete an address, select the address and click Delete.
Advanced	Specify the following information:
	• Cache —Select the check box to enable DNS caching and specify the following information:
	 Size—Specify the number of entries that the cache will hold (range 1024-10240, default 1024).
	 Timeout—Specify the length of time (hours) after which all cached entries are removed. DNS time-to-live values are used to remove cache entries when they have been stored for less than the configured timeout period. Following a timeout, new requests must be resolved and cached again (range 4 to 24, default 4 hours).
	 TCP Queries—Select the check box to enable DNS queries using TCP and specify the following information:
	 Max Pending Requests—Specify the upper limit on the number of concurrent pending TCP DNS requests that the firewall will support (range 64-256, default 64).
	• UDP Queries Retries—Specify settings for UDP query retries:
	 Interval—Specify the time in seconds after which another request is sent if no response has been received (range 1-30, default 2 seconds).
	 Attempts—Specify the maximum number of attempts (excluding the first attempt) after which the next DNS server is tried (range 1-30, default 5).

Table 72. DNS Proxy Settings (Continued)

Network Profiles

Network profiles capture configuration information that the firewall can use to establish network connections and implement policies. The following types of network profiles are supported:

- IKE gateways, IPSec crypto profiles, and IKE crypto profiles—These profiles support configuration and operation of IPSec VPNs. For information on the following profile types, refer to "Configuring IPSec Tunnels" on page 309.
 - IKE gateways include the configuration information that is necessary to perform IKE protocol negotiation with peer gateways when setting up IPSec VPN tunnels.
 - IKE crypto profiles specify the protocols and algorithms for Phase 1 identification, authentication, and encryption in VPN tunnels.
 - IPSec crypto profiles specify the protocols and algorithms for Phase 2 identification, authentication, and encryption in VPN tunnels.
- **Monitor profiles**—These profiles are used to monitor IPSec tunnels and to monitor a next-hop device for policy based forwarding (PBF) rules. In both cases, the monitor profile is used to specify an action to take when a resource (IPSec tunnel or next-hop device) becomes unavailable. Refer to "Monitor Settings" on page 177.
- Interface management profiles—These profiles specify the protocols that can be used to manage the firewall for Layer 3 interfaces, including VLAN and loopback interfaces. Refer to "Defining Interface Management Profiles" on page 175.
- **Zone protection profiles**—These profiles determine how the firewall responds to attacks from individual security zones. Refer to "Defining Zone Protection Profiles" on page 178. The following types of protection are supported:
 - **Flood Protection**—Protects against SYN, ICMP, UDP, and other IP-based flooding attacks.
 - **Reconnaissance detection**—Allows you to detect and block commonly used port scans and IP address sweeps that attackers run to find potential attack targets.
 - Packet-based attack protection—Protects against large ICMP packets and ICMP fragment attacks as well as a number of IP and TCP level attacks. Per-zone non-syn-tcp behavior can also be specified.
 - QoS profiles—These profiles determine how the QoS traffic classes are treated. You can set overall limits on bandwidth regardless of class and also set limits for individual classes. You can also assign priorities to different classes. Priorities determine how traffic is treated in the presence of contention. Refer to "Defining QoS Profiles" on page 358.

Defining Interface Management Profiles

Network > Network Profiles > Interface Mgmt

Use this page to specify the protocols that are used to manage the firewall. To assign management profiles to each interface, refer to "Configuring Layer 3 Interfaces" on page 130 and "Configuring Layer 3 Subinterfaces" on page 134. For an overview of firewall interfaces, refer to "Firewall Interfaces" on page 127.

Note: Although the default port used for HTTPS access is 443, if you enable HTTPS access on an inband port and Global Protect is enabled on the same port, the port used for HTTPS access is automatically set to 4443 instead of 443.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of interface management profiles when configuring interfaces. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Ping Telnet	Select the check box for each service to be enabled on the interfaces where the profile is applied.
SSH HTTP HTTPS	The Response Pages check box controls whether the ports used to serve captive portal and URL filtering response pages are open on Layer 3 interfaces. Ports 6080 and 6081 are left open if this setting is enabled.
SNMP Response Pages User-ID Service	The User-ID Service option is needed to allow communication between firewalls when a firewall is acting as a redistribution point to provide user mapping information to other PAN-OS firewalls. Refer to "User-ID
	Agent" on page 283.
Permitted IP Addresses	Enter the list of IPv4 or IPv6 addresses from which firewall management is allowed.

Table 73. Interface Management Profile Settings

Defining Monitor Profiles

Network > Network Profiles > Monitor

A monitor profile is used to monitor IPSec tunnels and to monitor a next-hop device for policy-based forwarding (PBF) rules. In both cases, the monitor profile is used to specify an action to take when a resource (IPSec tunnel or next-hop device) becomes unavailable. Monitor profiles are optional, but can be very useful for maintaining connectivity between sites and to ensure that PBF rules are maintained.

After creating a monitor profile, you can apply it to an IPSec Tunnel profile from the **General** tab, then selecting the **Show Advanced Options** check box, click the **Tunnel Monitor** check box and then select the desired profile from the drop-down list. In a PBF rule, click the **Forwarding** tab, click the **Monitor** check box and then select the desired profile from the drop down list.

In the example of an IPSec tunnel, the firewall monitors the specified IP address through the tunnel to determine if the tunnel is working properly according to the defined monitor profile. If the tunnel monitor IP address is not reachable, the action will be taken based on the settings you choose that are defined in "Monitor Settings" on page 177.

For a monitor profiles used in PBF rule, a remote IP address is monitored. If the remote IP address becomes unavailable, one of the following actions is taken.

- If the action is "wait-recover," packets continue to be sent according to the PBF rule.
- If the action is "fail-over," the firewall uses routing table lookup to determine routing for the duration of this session.

Field	Description
Name	Enter a name to identify the monitor profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Action	Specify an action to take if the tunnel is not available. If the threshold number of heartbeats is lost, the firewall takes the specified action.
	 wait-recover—Wait for the tunnel to recover; do not take additional action. Packets will continue to be sent according to the PBF rule.
	• fail-over —Traffic will fail over to a backup path, if one is available. The firewall uses routing table lookup to determine routing for the duration of this session.
	In both cases, the firewall tries to negotiate new IPSec keys to accelerate the recovery.
Interval	Specify the time between heartbeats (range 2-10, default 3).
Threshold	Specify the number of heartbeats to be lost before the firewall takes the specified action (range 2-100, default 5).

Table 74. Monitor Settings

Defining Zone Protection Profiles

▶ Network > Network Profiles > Zone Protection

Use this page to determine how the firewall responds to attacks from specified security zones. The same profile can be assigned to multiple zones. For an overview of security zones, refer to "Security Zones" on page 151. Flood protection applies to those packets that do not belong to an existing session.



Note: When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of zone protection profiles when configuring zones. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, and underscores.
Description	Enter an optional description for the zone protection profile.
Flood Protection Thresh	nolds - SYN Flood
Action	Select the action to take in response to a SYN flood attack.
	• Random Early Drop —Causes SYN packets to be dropped to mitigate a flood attack:
	 When the flow exceeds the Alert rate threshold, an alarm is generated.
	 When the flow exceeds the Activate rate threshold, individual SYN packets are dropped randomly to restrict the flow.
	 When the flow exceeds the Maximal rate threshold, all packets are dropped.
	• SYN Cookies —Computes a sequence number for SYN-ACK packets that does not require pending connections to be stored in memory. This is the preferred method.
Alert (packets/sec)	Enter the number of SYN packets received by the zone (in a second) that triggers an attack alarm. Alarms can be viewed on the Dashboard (refer to "Using the Dashboard" on page 252) and in the threat log (refer to "Identifying Unknown Applications and Taking Action" on page 276).
Activate (packets/sec)	Enter the number of SYN packets received by the zone (in a second) that triggers the action specified.
Maximum (packets/sec)	Enter the maximum number of SYN packets able to be received per second. Any number of packets exceeding the maximum will be dropped.
Flood Protection Thresh	nolds - ICMP Flood
Alert (packets/sec)	Enter the number of ICMP echo requests (pings) received per second that triggers an attack alarm.
Activate (packets/sec)	Enter the number of ICMP packets received by the zone (in a second) that causes subsequent ICMP packets to be dropped.

Table 75. Zone Protection Profile Settings

Field	Description
Maximum (packets/sec)	Enter the maximum number of ICMP packets able to be received per second. Any number of packets exceeding the maximum will be dropped.
Flood Protection Thresh	nolds - ICMPv6
Alert (packets/sec)	Enter the number of ICMPv6 echo requests (pings) received per second that triggers an attack alarm.
Activate (packets/sec)	Enter the number of ICMPv6 packets received per second for the zone that causes subsequent ICMPv6 packets to be dropped. Metering stops when the number of ICMPv6 packets drops below the threshold
Maximum (packets/sec)	Enter the maximum number of ICMPv6 packets able to be received per second. Any number of packets exceeding the maximum will be dropped.
Flood Protection Thresh	nolds - UDP
Alert (packets/sec)	Enter the number of UDP packets received by the zone (in a second) that triggers an attack alarm.
Alert (packets/sec)	Enter the number of UDP packets received by the zone (in a second) that triggers random dropping of UDP packets. The response is disabled when the number of UDP packets drops below the threshold.
Maximum (packets/sec)	Enter the maximum number of UDP packets able to be received per second. Any number of packets exceeding the maximum will be dropped.
Flood Protection Thres	nolds -Other IP
Alert (packets/sec)	Enter the number of IP packets received by the zone (in a second) that triggers an attack alarm.
Activate (packets/sec)	Enter the number of IP packets received by the zone (in a second) that triggers random dropping of IP packets. The response is disabled when the number of IP packets drops below the threshold. Any number of packets exceeding the maximum will be dropped.
Maximum (packets/sec)	Enter the maximum number of IP packets able to be received per second. Any number of packets exceeding the maximum will be dropped.
Reconnaissance Protec	tion - TCP Port Scan, UDP Port Scan, Host Sweep
Interval (sec)	Enter the time interval for port scans and host sweep detection (seconds).
Threshold (events)	Enter the number of scanned ports within the specified time interval that will trigger this protection type (events).
Action	Enter the action that the system will take in response to this event type:
	• Allow—Permits the port scan of host sweep reconnaissance.
	• Alert—Generates an alert for each scan or sweep that matches the threshold within the specified time interval.
	• Block —Drops all further packets from the source to the destination for the remainder of the specified time interval.
	• Block IP —Drops all further packets for a specified period of time. Choose whether to block source, destination, or source-and-destination traffic and enter a duration (seconds).

Table 75. Zone Protection Profile Settings (Continued)

Field	Description
IPv6 Drop Packets with	
Type 0 Router Header	Select the check box to drop IPv6 packets that include a Type 0 router header.
IPv4 Compatible Address	Select the check box to drop IPv6 packets that include an IPv4-compatible address.
Multicast Source Address	Select the check box to drop IPv6 packets that include a multicast source address.
Anycast Source Address	Select the check box to drop IPv6 packets that include an anycast source address.
Packet-Based Attack Pr	otection
TCP/IP Drop sub tab	
Spoofed IP address	Select the check box to enable protection against IP address spoofing.
Fragmented traffic	Discards fragmented IP packets.
Mismatched overlapping TCP	This setting will cause the firewall to report an overlap mismatch and drop the packet when segment data does not match in these scenarios:
segment	• The segment is within another segment.
	• The segment overlaps with part of another segment.
	• The segment covers another segment.
	This protection mechanism uses sequence numbers to determine where packets reside within the TCP data stream.
Remove TCP Timestamp	Determines whether the packet has a TCP timestamp in the header and, if it does, strips the timestamp from the header.
Reject Non-SYN TCP	Determines whether to reject the packet, if the first packet for the TCP session setup is not a SYN packet:
	• global —Use system-wide setting that is assigned through the CLI.
	• yes —Reject non-SYN TCP.
	• no —Accept non-SYN TCP. Note that allowing non-SYN TCP traffic may prevent file blocking policies from working as expected in cases where the client and/or server connection is not set after the block occurs.
Asymmetric Path	Determine whether to drop or bypass packets that contain out of sync ACKs or out of window sequence numbers:
	• global —Use system wide setting that is assigned through the CLI.
	• drop —Drop packets that contain an asymmetric path.
	• bypass —Bypass scanning on packets that contain an asymmetric path.
IP Option Drop	
Strict Source Routing	Discard packets with the Strict Source Routing IP option set.
Loose Source Routing	Discard packets with the Loose Source Routing IP option set.
Timestamp	Discard packets with the Timestamp IP option set.
Record Route	Discard packets with the Record Route IP option set.
Security	Discard packets if the security option is defined.

Table 75. Zone Protection Profile Settings (Continue
--
Field
--
Stream ID
Unknown
Malformed
ICMP Drop sub tab
ICMP ping ID 0
ICMP fragment
ICMP large packet (>1024)
Suppress ICMP TTL expired error
Suppress ICMP Frag Needed
IPv6 Drop sub tab
Type 0 Routing Header
IPv4 compatible address
Anycast source address
Needless fragment header
MTU in ICMPv6 'Packet Too Big' less than 1280 bytes
Hop-by-Hop extension
Routing extension
Destination extension
Invalid IPv6 options in extension header
Non-zero reserved field
ICMPv6 sub tab
ICMPv6 destination unreachable - require explicit security rule match
ICMPv6 packet too big - require explicit security rule match
ICMPv6 time exceeded - require explicit security rule match

Table 75. Zone Protection Profile Settings (Continued)

Field	Description
ICMPv6 parameter problem - require explicit security rule match	Require an explicit security policy match for parameter problem ICMPv6 errors even when associated with an existing session.
ICMPv6 redirect - require explicit security rule match	Require an explicit security policy match for redirect ICMPv6 messages even when associated with an existing session.

 Table 75.
 Zone Protection Profile Settings (Continued)

Chapter 5 Policies and Security Profiles

This chapter describes how to configure policies and security profiles:

- "Policies" in the next section
- "Security Profiles" on page 211
- "Other Policy Objects" on page 226

Policies

Policies allow you to control firewall operation by enforcing rules and automatically taking action. The following types of policies are supported:

- Basic security policies to block or allow a network session based on the application, the source and destination zones and addresses, and optionally the service (port and protocol). Zones identify the physical or logical interfaces that send or receive the traffic. Refer to "Security Policies" on page 187.
- Network Address Translation (NAT) policies to translate addresses and ports, as needed. Refer to "NAT Policies" on page 190.
- Policy-based forwarding policies to determine the egress interface used following processing. Refer to "Policy-Based Forwarding Policies" on page 199.
- Decryption policies to specify traffic decryption for security policies. Each policy can specify the categories of URLs for the traffic you want to decrypt. SSH decryption is used to identify and control SSH tunneling in addition to SSH shell access. Refer to "Decryption Policies" on page 202.
- Override policies to override the application definitions provided by the firewall. Refer to "Application Override Policies" on page 205.
- Quality of Service (QoS) policies to determine how traffic is classified for treatment when it passes through an interface with QoS enabled. Refer to "Defining QoS Policies" on page 359.

- Captive portal policies to request authentication of unidentified users. Refer to "Captive Portal Policies" on page 206.
- Denial of service (DoS) policies to protect against DoS attacks and take protective action in response to rule matches. Refer to "DoS Protection Policies" on page 209.



Note: Shared polices pushed from Panorama are shown in green on the firewall web interface pages and cannot be edited at the device level.

Guidelines on Defining Policies

For general guidelines on interacting with the firewall interface, refer to "Using the Firewall Web Interface" on page 23. The following specific guidelines apply when interacting with the pages on the **Policies** tab:

- To apply a filter to the list, select from the **Filter Rules** drop-down list. To add a value to define a filter, click the down-facing arrow for the item and choose **Filter**.
- To view application groups, filters, or container information when creating or viewing Security, PBF, or QoS policies, hold your mouse over the object in the **Application** column, click the down arrow and select **Value**. This allows you to easily view application members directly from the policy without having to navigate to the Object tabs.
- To add a new policy rule, do one of the following:
 - Click **Add** at the bottom of the page.
 - Select a rule on which to base the new rule and click Clone Rule, or select a rule by clicking the white space of the rule, and select Clone Rule at the bottom of the page (a selected rule has a yellow background). The copied rule, "rulen" is inserted below the selected rule, where *n* is the next available integer that makes the rule name unique.
- The order in which rules are listed is the order in which the rules are compared against network traffic. Change the ordering of a rule in either of the following ways:
 - Select the rule and click Move Up, Move Down, Move Top, or Move Bottom.
 - Click the down-facing arrow for the rule name and choose **Move**. In the pop-up window, choose a rule and choose whether to move the rule you selected for reordering before or after this rule.

Block MSN	Block MSN	😤 Filter	>	Move - Block MSN		0
	i.	Log Viewer	t	Select a Rule		•
		Move	t		Move Before SMove After	Cancel
		Lec (h) any				

• To enable a rule, select the rule and click **Enable**.

• To show which rules are not currently used, select the **Highlight Unused Rules** check box.

Name	Tag	Zone	Address	User	HIP Profile	Zone	Address	Application	Service	Action
Dợ Nơi Log Traffic	No Log filber	(RA), tapzone	đely	ahy	ariy.	(M) Capzone	😭 LocalServers	đny	a02	ø
Do Not Log URL	No Log	🚧 tapzone	any	any	any	🕅 tapzone	Se LocalNetwork	🔝 ssl	any	0
▲								📰 web-browsing		
-										
ule used			Rule not us	ed (yellov	v dotted k	ackgroun	d)			

- To display the log for the policy, click the down-facing arrow for the rule name and choose **Log Viewer**.
- For some entries, you can display the current value by clicking the down-facing arrow for the entry and choosing **Value**. You can also edit, filter, or remove certain items directly from the column menu.

Sou							
Address	User		HIP Profile	Zone	Address		
Sig Corp-Mktg	Edit		any	🙀 13-untrust	any		
G	Filter						
	Remove						
	Value 🕨 🕨		Addr	ess			
		Description: Test					
		IP Netmask: 192.168.1.5					

• If you have a large number of policies defined, you can use the filter bar to find objects that are used within a policy based on the object name or IP address. The search will also traverse embedded objects to find an address within an address object or address group. In the following screen shot, the IP address 10.8.10.177 was entered in the filter bar and the policy "aaa" is shown. That policy uses an address group object named "aaagroup", which contains the IP address.

m naloalto		Filter bar								
NETWORKS	Dashboar	d	ACC	Menitor	Policies	Objects	Network	Device		
	Virtua	al System vs y	ys1		~					
Security NAT Cos		10.8.10.177					Source			
Policy Based Forwarding		Name	Tag	Zone	Address			User	HIP Profile	Zone
Decryption Application Override	1	aaa	none	any	🚱 aaag n 🥵 bbbgrou	oup Ip		any 💌	any	any
DoS Protection			F	ilter res	sults					

• You can show or hide specific columns from view in any of the **Policies** pages.



Specifying Users and Applications for Policies

- Policies > Security
- Policies > Decryption

You can restrict security policies to selected users or applications by clicking the **User** or **Application** link on the **Security** or **Decryption** device rules page. For information on restricting rules by application, refer to "Defining Applications" on page 233.

To restrict a policy to selected users, follow these steps:

1. On the **Security** or **Decryption** device rules page, click the underlined link for the source or destination user to open the selection window.



Note: If you are using a RADIUS server and not the User-ID Agent, the list of users is not displayed, and you must enter user information manually.

- 2. Choose the type of rule to apply:
 - **any**—Includes any user in the rule.
 - pre-logon—Include remote users that are connected to the network using GlobalProtect, but are not logged into their system. When the Pre-logon option is configured on the Portal for GlobalProtect clients, any user who is not currently logged into their machine will be identified with the username pre-logon. You can then create policies for pre-logon users and the machines can be accessed as if they were logged in and connected to the domain.
 - known-user—Includes all authenticated users.
 - **unknown**—Includes all unauthenticated users.
 - Select—Includes selected users as determined by the selection in this window.
- 3. To add groups of users, select from the **Available User Groups** check boxes and click **Add User Group**. Alternatively, you can enter text to match one or more groups and click **Add User Group**.

- 4. To add individual users, enter a search string in the **User** search field and click **Find**. You can then select users and click **Add User**. Alternatively, you can enter individual user names in the **Additional Users** area.
- 5. Click **OK** to save the selections and update the security or decryption rule.

Security Policies

Security policies determine whether to block or allow a new network session based on traffic attributes such as the application, source and destination security zones, the source and destination addresses, and the application service (such as HTTP). Security zones are used to group interfaces according to the relative risk of the traffic they carry. For example, an interface connected to the Internet is in an "untrusted" zone, while an interface connected to the internal network is in a "trusted" zone.



Note: By default, traffic between each pair of security zones is blocked until at least one rule is added to allow traffic between the two zones.

Intra-zone traffic is allowed by default and requires an explicit block rule. If a deny all rule is added as the last rule in the policy, intra-zone traffic will be blocked unless otherwise allowed.

Security policies can be as general or specific as needed. The policy rules are compared against the incoming traffic in sequence, and because the first rule that matches the traffic is applied, the more specific rules must precede the more general ones. For example, a rule for a single application must precede a rule for all applications if all other traffic-related settings are the same.

Defining Security Policies

Policies > Security

Use the **Security** page to define security policy rules. For configuration guidelines, refer to "Guidelines on Defining Policies" on page 184.

Field	Description
General Tab	
Name	Enter a name to identify the rule (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required.
Description	Enter a description for the policy (up to 255 characters).
Tag	If you need to tag the policy, click Add to specify the tag. A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location.

Table 76. Security Policy Settings

Field	Description
Source Tab	
Source Zone	Click Add to choose source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.
	Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases.
Source Address	Click Add to add source addresses, address groups, or regions (default is any). Select from the drop-down list, or click the Address , Address Group , or Regions link at the bottom of the drop-down list, and specify the settings.
User Tab	
Source User	Click Add to choose the source users or groups of users subject to the policy.
HIP Profiles	Click Add to choose Host Information Profiles (HIP) to identify users. For information on HIP, refer to "Overview" on page 335.
Destination Tab	
Destination Zone	Click Add to choose destination zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.
	Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases.
Destination Address	Click Add to add destination addresses, address groups, or regions (default is any). Select from the drop-down list, or click the Address link at the bottom of the drop-down list, and specify address settings.
Application Tab	
Application	Select specific applications for the security rule. To define new applications, refer to "Defining Applications" on page 233. To define application groups, refer to "Defining Application Groups" on page 238.
	If an application has multiple functions, you can select the overall application or individual functions. If you select the overall application, all functions are included, and the application definition is automatically updated as future functions are added.
	If you are using application groups, filters, or container in the security rule, you can view details on these objects by holding your mouse over the object in the Application column, click the down arrow and select Value . This allows you to easily view application members directly from the policy without having to navigate to the Object tabs.

Table 76. Security Policy Settings (Continued)

Field	Description						
Service/ URL Category Tab							
Service	Select services to limit to specific TCP and/or UDP port numbers. Choose one of the following from the drop-down list:						
	 any—The selected applications are allowed or denied on any protocol or port. 						
	• application-default —The selected applications are allowed or denied only on their default ports defined by Palo Alto Networks. This option is recommended for allow policies because it prevents applications from running on unusual ports and protocols, which if not intentional, can be a sign of undesired application behavior and usage. Note that when you use this option, the device still checks for all applications on all ports, but with this configuration, applications are only allowed on their default ports/protocols.						
	• Select—Click Add. Choose an existing service or choose Service or Service Group to specify a new entry. Refer to "Services" on page 239 and "Service Groups" on page 240.						
URL Category	Select URL categories for the security rule.						
	• Choose any to allow or deny all sessions regardless of the URL category.						
	• To specify a category, click Add and select a specific category (including a custom category) from the drop-down list. You can add multiple categories. Refer to "Custom URL Categories" on page 242 for information on defining custom categories.						
Actions Tab							
Action Setting	Click allow or deny to allow or block a new network session for traffic that matches this rule.						
Profile Setting	To specify the checking done by the default security profiles, select individual Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking, and/or Data Filtering profiles.						
	To specify a profile group, rather than individual profiles, select Profile Type Group and then select a profile group from the Group Profile drop- down list.						
	To define new profiles or profile groups, click New next to the appropriate profile or group (refer to "Security Profile Groups" on page 246).						

Table 76. Security Policy Settings (Continued)

Field	Description
Log Setting	Specify any combination of the following options:
	Log Setting:
	• To forward the local traffic log and threat log entries to remote destina- tions, such as Panorama and syslog servers, select a log profile from the Log Forwarding Profile drop-down list. Note that the generation of threat log entries is determined by the security profiles. To define new log profiles, click New (refer to "Log Forwarding" on page 247).
	• To generate entries in the local traffic log for traffic that matches this rule, select the following options:
	 Log At Session Start. Generates a traffic log entry for the start of a session (disabled by default).
	 Log At Session End. Generates a traffic log entry for the end of a session (enabled by default).
	If the session start or end entries are logged, drop and deny entries are also logged.
Other Settings	Specify any combination of the following options:
	• Schedule—To limit the days and times when the rule is in effect, select a schedule from the drop-down list. To define new schedules, click New (refer to "Schedules" on page 250).
	• QoS Marking —To change the Quality of Service (QoS) setting on packets matching the rule, select IP DSCP or IP Precedence and enter the QoS value in binary or select a predefined value from the drop-down list. For more information on QoS, refer to "Configuring Quality of Service" on page 355.
	• Disable Server Response Inspection —To disable packet inspection from the server to the client, select this check box. This option may be useful under heavy server load conditions.

Table 76. Security Policy Settings (Continued)

NAT Policies

If you define Layer 3 interfaces on the firewall, you can use Network Address Translation (NAT) policies to specify whether source or destination IP addresses and ports are converted between public and private addresses and ports. For example, private source addresses can be translated to public addresses on traffic sent from an internal (trusted) zone to a public (untrusted) zone.

NAT is also supported on virtual wire interfaces. When performing NAT on virtual wire interfaces, it is recommended that you translate the source address to a different subnet than the one on which the neighboring devices are communicating. Proxy ARP is not supported on virtual wires and so neighboring devices will only be able to resolve ARP requests for IP addresses that reside on the interface of the device on the other end of the virtual wire.

When configuring NAT on the firewall, it is important to note that a security policy must also be configured to allow the NAT traffic. Security policy will be matched based on the post-NAT zone and the pre-NAT IP address.

The firewall supports the following types of address translation:

• **Dynamic IP/Port**—For outbound traffic. Multiple clients can use the same public IP addresses with different source port numbers. Dynamic IP/Port NAT rules allow translation to a single IP address, a range of IP addresses, a subnet, or a combination of

these. In cases where an egress interface has a dynamically assigned IP address, it can be helpful to specify the interface itself as the translated address. By specifying the interface in the dynamic IP/port rule, NAT policy will update automatically to use any address acquired by the interface for subsequent translations.



Note: Palo Alto Networks Dynamic IP/port NAT supports more NAT sessions than are supported by the number of available IP addresses and ports. The firewall can use IP address and port combinations up to two times (simultaneously) on the PA-200, PA-500, PA-2000 Series and PA-3000 Series, four times on the PA-4020 and PA-5020, and eight times on the PA-4050, PA-4060, PA-5050, and PA-5060 devices when destination IP addresses are unique.

- **Dynamic IP**—For outbound traffic. Private source addresses translate to the next available address in the specified address range. Dynamic IP NAT policies allow you to specify a single IP address, multiple IPs, multiple IP ranges, or multiple subnets as the translated address pool. If the source address pool is larger than the translated address pool, new IP addresses seeking translation will be blocked while the translated address pool is fully utilized. To avoid this issue, you can specify a fall back pool that will be used if the primary pool runs out of IP addresses.
- **Static IP**—For inbound or outbound traffic. You can use static IP to change the source or the destination IP address while leaving the source or destination port unchanged. When used to map a single public IP address to multiple private servers and services, destination ports can stay the same or be directed to different destination ports.



Note: You may need to define static routes on the adjacent router and/or the firewall to ensure that traffic sent to a public IP address is routed to the appropriate private address. If the public address is the same as the firewall interface (or on the same subnet), then a static route is not required on the router for that address. When you specify service (TCP or UDP) ports for NAT, the pre-defined HTTP service (service-http) includes two TCP ports: 80 and 8080. To specify a single port, such as TCP 80, you must define a new service.

The next table summarizes the NAT types. The two dynamic methods map a range of client addresses (M) to a pool (N) of NAT addresses, where M and N are different numbers. N can also be 1. Dynamic IP/Port NAT differs from Dynamic IP NAT in that the TCP and UDP source ports are not preserved in Dynamic IP/Port, whereas they are unchanged with Dynamic IP NAT. There are also differing limits to the size of the translated IP pool, as noted below.

With Static IP NAT, there is a one-to-one mapping between each original address and its translated address. This can be expressed as 1-to-1 for a single mapped IP address, or M-to-M for a pool of many one-to-one, mapped IP addresses.

PAN-OS NAT Type	Source Port Stays the Same	Destination Port Can Change	Mapping Type	Size of Translated Address Pool
Dynamic IP/ Port	No	No	Many-to-1 M-to-N	Up to 254 consecutive addresses
Dynamic IP	Yes	No	M-to-N	Up to 32k consecutive addresses
Static IP	Yes	No	1-to-1 M-to-M MIP	Unlimited
	Optional		1-to-Many VIP PAT	

Table 77. NAT Types

Determining Zone Configuration in NAT and Security Policy

NAT rules must be configured to use the zones associated with pre-NAT IP addresses configured in the policy. For example, if you are translating traffic that is incoming to an internal server (which is reached via a public IP by Internet users), it is necessary to configure the NAT policy using the zone in which the public IP address resides. In this case, the source and destination zones would be the same. As another example, when translating outgoing host traffic to a public IP address, it is necessary to configure NAT policy with a source zone corresponding to the private IP addresses of those hosts. The pre-NAT zone is required because this match occurs before the packet has been modified by NAT.

Security policy differs from NAT policy in that post-NAT zones must be used to control traffic. NAT may influence the source or destination IP addresses and can potentially modify the outgoing interface and zone. When creating security policies with specific IP addresses, it is important to note that pre-NAT IP addresses will be used in the policy match. Traffic subject to NAT must be explicitly permitted by the security policy when that traffic traverses multiple zones.

NAT Rule Options

The firewall supports no-NAT rules and bi-directional NAT rules.

No-NAT Rules

No-NAT rules are configured to allow exclusion of IP addresses defined within the range of NAT rules defined later in the NAT policy. To define a no-NAT policy, specify all of the match criteria and select **No Source Translation** in the source translation column.

Bi-directional NAT Rules

The bi-directional setting in static source NAT rules implicitly creates a destination NAT rule for traffic to the same resources in the reverse direction. In this example, two NAT rules are used to create a source translation for outgoing traffic from IP 10.0.1.10 to public IP 3.3.3.1 and a destination translation for traffic destined for public IP 3.3.3.1 to private IP 10.0.1.10. This pair of rules can be simplified by configuring only the third NAT rule using the bi-directional feature.

Name	Tag	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
sourcenat	none	🙀 L3Trust	M L3Untrust	any	5 10.0.1.10	any	any	static-ip	none
								3.3.3.1	
								bi-directional: false	
destnat	none	M L3Untrust	Magazine La Trust	any	any	5.3.3.1	any	none	address: 10.0.1.10
bothnat	none	M L3Trust	Magazing L3Untrust	any	5 10.0.1.10	any	any	static-ip	none
								3.3.3.1	
								bi-directional: true	

Figure 17. Bi-Directional NAT Rules

Defining Network Address Translation Policies

Policies > NAT

NAT address translation rules are based on the source and destination zones, the source and destination addresses, and the application service (such as HTTP). Like security policies, the NAT policy rules are compared against the incoming traffic in sequence, and the first rule that matches the traffic is applied.

As needed, add static routes to the local router so that traffic to all public addresses is routed to the firewall. You may also need to add static routes to the receiving interface on the firewall to route traffic back to the private address (refer to "Firewall Interfaces" on page 127). For configuration guidelines, refer to "Guidelines on Defining Policies" on page 184.

Field	Description
Name	Change the default rule name and/or enter a rule description.
Description	Enter a description for the policy (up to 255 characters).
NAT Type	Specify ipv4 for NAT between IPv4 addresses, or nat64 translation between IPv6 and IPv4 addresses.
	<i>Note:</i> You cannot combine IPv4 and IPv6 address ranges in a single NAT rule.
Tag	If you need to tag the policy, click Add to specify the tag.
	A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location.
Original Packet	
Source Zone Destination Zone	Select one or more source and destination zones for the original (non-NAT) packet (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.
	Multiple zones can be used to simplify management. For example, you can configure settings so that multiple internal NAT addresses are directed to the same external IP address.
Destination Interface	Specify the type of interface for translation. The destination interface can be used to translate IP addresses differently in the case where the network is connected to two ISPs with different IP address pools.
Source Address Destination Address	Specify a combination of source and destination addresses for which the source or destination address must be translated.
Service	Specify the services for which the source or destination address is translated. To define new service groups, refer to "Service Groups" on page 240.

Table 78. NAT Rule Settings

Field	Description		
Translated Packet			
Source Translation	Enter an IP address or address range (address1-address2) that the source address is translated to, and select a dynamic or static address pool. The size of the address range is limited by the type of address pool:		
	• Dynamic IP And Port—Address selection is based on a hash of the source IP address. For a given source IP address, the firewall will use the same translated source address for all sessions. Dynamic IP and Port source NAT supports approximately 64k concurrent sessions on each IP address in the NAT pool. On some platforms, over-subscription is supported, which will allow a single IP to host more than 64k concurrent sessions. See the note on page 191 for more details.		
	• Dynamic IP —The next available address in the specified range is used, but the port number is unchanged. Up to 32k consecutive IP addresses are supported. A dynamic IP pool can contain multiple subnets, so you can translate your internal network addresses to two or more separate public subnets.		
	- Advanced (Fall back Dynamic IP Translation)—Use this option to create a fall back pool that will perform IP and port translation and will be used if the primary pool runs out of addresses. You can define addresses for the pool by using the Translated Address option or the Interface Address option, which is for interfaces that receive an IP address dynamically. When creating a fall back pool, make sure addresses do not overlap with addresses in the primary pool.		
	• Static IP —The same address is always used, and the port is unchanged. For example, if the source range is 192.168.0.1-192.168.0.10 and the translation range is 10.0.0.1-10.0.0.10, address 192.168.0.2 is always translated to 10.0.0.2. The address range is virtually unlimited.		
	• None—Translation is not performed.		
Destination Translation	Enter an IP address or range of IP addresses and a translated port number (1 to 65535) that the destination address and port number are translated to. If the Translated Port field is blank, the destination port is not changed. Destination translation is typically used to allow an internal server, such as an email server, to be accessed from the public network.		

Table 78. NAT Rule Settings (Continued)

NAT Policy Examples

The following NAT policy rule translates a range of private source addresses (10.0.0.1 to 10.0.0.100 in the "L3Trust" zone) to a single public IP address (200.10.2.100 in the "L3Untrust" zone) and a unique source port number (dynamic source translation). The rule applies only to traffic received on a Layer 3 interface in the "L3Trust" zone that is destined for an interface in the "L3Untrust" zone. Because the private addresses are hidden, network sessions cannot be initiated from the public network. If the public address is not a firewall interface address (or on the same subnet), the local router requires a static route to direct return traffic to the firewall.

Security policy must be explicitly configured to permit traffic matching this NAT rule. Create a security policy with source/destination zones and source/destination addresses matching the NAT rule.

			Original Packet					Transi	ated Packet
Name	Tag	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
Client Source NAT	none	🕅 L3Trust	pag L3Untrust	any	5 10.0.0.1-10.0.0.100	any	any	dynamic-ip-and-port	none

Figure 18. Dynamic Source Address Translation

In the following example, the first NAT rule translates the private address of an internal mail server to a static public IP address. The rule applies only to outgoing email sent from the "L3Trust" zone to the "L3Untrust" zone. For traffic in the reverse direction (incoming email), the second rule translates the destination address from the server's public address to its private address. Rule2 uses "L3Untrust" for the source and destination zones because NAT policy is based on the pre-NAT address zone. In this case, that pre-NAT address is a public IP address and is therefore in the "L3Untrust" zone.

Name	Tag	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
rule1	none	Ma L3Trust	M L3Untrust	any	🔄 Private Email	any	any	static-ip 200.10.2.100 bi-directional: false	none
rule2	none	M L3Untrust	pm L3Untrust	any	any	Seg Public Email	any	none	address: 192.168.2.200

Figure 19. Static Source and Destination Address Translation

In both examples, if the public address is not the address of the firewall's interface (or on the same subnet), you must add a static route to the local router to route traffic to the firewall.

NAT64

NAT64 is used to translate source and destination IP headers between IPv6 and IPv4 addresses. It allows IPv6 clients to access IPv4 servers and allows IPv4 clients to access IPv6 servers. There are three main transition mechanisms defined by the IETF: dual-stack, tunneling, and translation. When you have IPv4-only and IPv6-only networks and communication is required, you must use translation.

When using NAT64 policies on Palo Alto Networks firewall, it's required that you have a third-party DNS64 solution in place to decouple the DNS query function from the NAT function.

The following NAT64 features are supported:

- Stateful NAT64, which allows preserving of IPv4 addresses so that one IPv4 address can map to multiple IPv6 addresses. An IPv4 address can also be shared with NAT44. In contrast, stateless NAT64 maps one IPv4 address to one IPv6 address.
- IPv4-initiated communication translation. The static binding for IPv4 maps an IPv4 address/port number to an IPv6 IP address. PAN-OS also supports port rewrite that allows you to preserve even more IPv4 addresses.
- Allows translation for /32, /40, /48, /56, /64, and /96 subnets.
- Support for multiple prefixes. You can assign one NAT64 prefix per rule.
- Does not require you to reserve a pool of IPv4 address specifically for NAT64. Therefore, it is possible to use a single IP address to do NAT44 and NAT64.
- Supports hairpinning (NAT U-Turn) and it can prevent hairpinning loop attacks.

- Supports the translation of TCP/UDP/ICMP packets as per RFC, but also other protocols without ALG (best effort). For example, a GRE packet could be translated. This translation has the same limitation as NAT44.
- Supports PMTUD (path MTU discovery) and it updates the MSS (Maximum Segment Size) for TCP.
- Allows configuration of the IPv6 MTU setting. The default value is 1280, which is the minimum MTU for IPv6 traffic. This setting is configured in **Device > Setup > Sessions** tab under **Session Settings**.
- Translates length attribute between IPv4 and IPv6.
- Supported on Layer 3 interfaces and subinterfaces, tunnel, and VLAN interfaces.

NAT64 Examples

You can configure two types of translation with the firewall: IPv6-initiated communication, which is similar to source NAT in IPv4, and IPv4-initiated communication to an IPv6 server, which is similar to destination NAT in IPv4.

IPv6-initiated Communication

In this type of translation, the destination IPv6 address in the NAT rule is a prefix following the RFC 6052 format (/32, /40,/48,/56,/64, and /96). The destination IPv6 address netmask in the rule would be used to extract the IPv4 address. The source translation needs to have "Dynamic IP and Port" in order to implement a stateful NAT64. The IPv4 address set as the source is configured the same way as a NAT44 source translation. The destination translation field is not set. However, a destination translation must be done since the address is extracted from the IPv6 address in the packet. It uses prefix defined in the destination IP matching criteria. You should note that in a /96 prefix, it is the last 4 octets, but the location of the IPv4 address would be different if the prefix is not /96.



Figure 20. NAT64 IPv6 Client to IPv4 Network Diagram

The following	table describes	s the values nee	eded in this N	AT64 policy.
0	/			1 /

Source IP	Destination IP	Source Translation	Destination Translation
Any/IPv6 address	NAT64 IPv6 prefix with RFC6052 compliant netmask	Dynamic IP and port mode (Use IPv4 address)	None (Extracted from the destination IPv6 address)

Table 79.

IPv4-initiated Communication

The IPv4 address is the address that maps to the IPv6 address and you use static IP mode in the source translation. The source would be set in an IPv6 prefix as defined in RFC6052 and is appended to the IPv4 source address. The destination address is the IP address set in the destination translation column. It is possible to rewrite the destination port. This method allows a single IP address to share multiple IPv6 servers through the port through a static mapping.



Figure 21. NAT64 IPv4 Internet to IPv6 Customer Network Diagram

The following table describes the values needed in this NAT64 policy.

Source IP	Destination IP	Source Translation	Destination Translation
Any/IPv4 address	IPv4 address	Static IP mode (IPv6 prefix in RFC	Single IPv6 address (actual server IP address)
		6052 format)	Note: You could specify a server port re-write.

Table 80. IPv4-initiated Values

The packet processing engine of the firewall must do a route lookup to find the destination zone prior to looking at the NAT rule. In NAT64, it is important to address the reachability of the NAT64 prefix for the destination zone assignment since the NAT64 prefix should not be routable by the NAT64 gateway. It is very likely that the NAT64 prefix would hit the default route, or be dropped because there is no route. You can setup a tunnel interface with no termination point since this type of interface will act like a loopback port and accept other netmasks besides /128. You apply the NAT64 prefix to the tunnel and apply the appropriate

zone in order to ensure that IPv6 traffic with NAT64 prefix is assigned to the proper destination zone. It would also have the advantage to drop the IPv6 traffic with NAT64 prefix if the NAT64 rule is not matched.

IETF Scenarios for IPv4/IPv6 Translation

There are six NAT64 based scenarios defined by the IETF in RFC 6144. The Palo Alto Networks firewall will support all but one of these scenarios, as described in the following table.

Scenario	Source IP	Destination IP	Source Translation	Destination Translation
IPv6 Network to the IPv4 Internet	Any/IPv6 address	NAT64 IPv6 prefix with RFC 6052 compliant netmask.	Dynamic IP and port mode. Use <i>Public</i> IPv4 address	None (extracted from destination IPv6 address)
The IPv4 Internet to an IPv6 Network	Any/IPv4 address	Single IPv4 address	Static IP mode. IPv6 Prefix in RFC 6052 format	Single IPv6 address
The IPv6 Internet to an IPv4 Network	Any/IPv6 address	IPV6 globally routable prefix with RFC 6052 compliant netmask.	Dynamic IP and port. Use <i>Private</i> IPv4 address	None (extracted from destination IPv6 address)
IPv4 network to IPv6 Internet	Not currently	supported		
IPv4 network to IPv6 network	Any/IPv4 address	Single IPv4 address	Static IP mode. IPv6 Prefix in RFC 6052 format	Single IPv6 address
IPv6 network to IPv4 network	Any/IPv6 address	NAT64 IPV6 prefix with RFC 6052 compliant netmask.	Dynamic IP and port. Use Private IPv4 address	None (extracted from destination IPv6 address)

 Table 81.
 Summary of IETF Scenario Implementations for Using PAN-OS

Policy-Based Forwarding Policies

Policies > Policy Based Forwarding

Normally, when traffic enters the firewall, the ingress interface virtual router dictates the route that determines the outgoing interface and destination security zone based on destination IP address. With policy-based forwarding (PBF), you can specify other information to determine the outgoing interface, including source zone, source address, source user, destination address, destination application, and destination service. The initial session on a given destination IP address and port that is associated with an application will not match an application-specific rule and will be forwarded according to subsequent PBF rules (that do not specify an application) or the virtual router's forwarding table. All subsequent sessions on

that destination IP address and port for the same application will match an applicationspecific rule. To ensure forwarding through PBF rules, application-specific rules are not recommended.

When necessary, PBF rules can be used to force traffic through an additional virtual system using the Forward-to-VSYS forwarding action. In this case, it is necessary to define an additional PBF rule that will forward the packet from the destination virtual system out through a particular egress interface on the firewall.

For configuration guidelines, refer to "Guidelines on Defining Policies" on page 184.

Field	Description
General Tab	
Name	Enter a name to identify the rule (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required.
Description	Enter a description for the policy (up to 255 characters).
Tag	If you need to tag the policy, click Add to specify the tag.
	A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location.
Source Tab	
Source Zone	To choose source zones (default is any), click Add and select from the drop-down list. To define new zones, refer to "Defining Security Zones" on page 151.
	Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases.
	Note: Only Layer 3 type zones are supported for policy-based forwarding.
Source Address	Click Add to add source addresses, address groups, or regions (default is any). Select from the drop-down list, or click the Address , Address Group , or Regions link at the bottom of the drop-down list, and specify the settings.
Source User	Click Add to choose the source users or groups of users subject to the policy.
Destination/ Application/Service Tab	
Destination Address	Click Add to add destination addresses, address groups, or regions (default is any). Select from the drop-down list, or click the Address , Address Group , or Regions link at the bottom of the drop-down list, and specify the settings.

Table 82. Policy-Based Forwarding Settings

Field	Description	
Application	Select specific applications for the PBF rule. To define new applications, refer to "Defining Applications" on page 233. To define application groups, refer to "Defining Application Groups" on page 238.	
	If you are using application groups, filters, or container in the PBF rule, you can view details on these objects by holding your mouse over the object in the Application column, click the down arrow and select Value . This enables you to easily view application members directly from the policy without having to go to the Object tabs.	
Service	Specify the services for which the source or destination address is translated. To define new service groups, refer to "Service Groups" on page 240.	
Forwarding Tab		
Action	Select one of the following options:	
	 Forward—Specify the next hop IP address and egress interface (the interface that the packet takes to get to the specified next hop). 	
	• Forward To VSYS—Choose the virtual system to forward to from the drop-down list.	
	• Discard —Drop the packet.	
	• No PBF —Do not alter the path that the packet will take.	
Egress Interface	Specify the firewall interface for forwarding traffic from the firewall.	
Next Hop	Specify the IP address of the next forwarding stop.	
Monitor	To monitor the forwarding actions, select Monitor and specify the following settings:	
	• Profile —Choose a profile from the drop-down list.	
	• Disable if unreachable —Select this check box if you want to ignore this rule for all new sessions when the next hop router is unreachable.	
	• IP Address —Specify the IP address to which ping messages are sent periodically to determine the state of the policy based forwarding rule.	

Table 82. Policy-Based Forwarding Settings (Continued)

Field	Description
Enforce Symmetric Return	Select this option to allow virtual routers on the firewall to enforce symmetric return for the PBF rule. Click Add if you want to enter the next-hop address(es) that will be used for the forwarding.
	This option will circumvent the route lookup process for return traffic; the firewall will use the original incoming interface as the egress interface. If the source IP is in the same subnet as the incoming interface on the firewall, symmetric return will not take effect.
	This feature is useful when you have servers accessible through two ISP connections (on different ingress interfaces) and the return traffic must be routed through the ISP that original routed the sessions.
	The following should be noted when using this feature:
	• If an interface has multiple PBF rules, only one rule can enforce symmetric return.
	• You can use this option for all L3 interfaces, except loopback interfaces. You can also use interfaces that have the IP address assigned dynami- cally (DHCP and PPoE).
	• The source must be an interface, not a zone.
	• Next-hop address list is not supported for tunnel and PPoE interfaces.
	• Up to 8 next-hop addresses can be defined per PBF rule.
Schedule	To limit the days and times when the rule is in effect, select a schedule from the drop-down list. To define new schedules, refer to "Schedules" on page 250.

Table 82. Policy-Based Forwarding Settings (Continued)

Decryption Policies

▶ Policies > Decryption

You can configure the firewall to decrypt traffic for visibility, control, and granular security. Decryption policies can apply to Secure Sockets Layer (SSL) and Secure Shell (SSH) traffic. With the SSH option, you can selectively decrypt outbound and inbound SSH traffic to assure that secure protocols are not being used to tunnel disallowed applications and content. You can also apply decryption profiles to your policies to block and control various aspects of SSL traffic. For more information, refer to "Decryption Profiles" on page 248.

Each decryption policy specifies the categories of URLs to decrypt or not decrypt. App-ID and the Antivirus, Vulnerability, Anti-spyware, URL Filtering, and File-blocking profiles are applied to decrypted traffic before it is re-encrypted as traffic exits the device. End-to-end security between clients and servers is maintained, and the firewall acts as a trusted third party during the connection. No decrypted traffic leaves the device.

The firewall inspects traffic, regardless of the protocols that are encapsulated. Decryption policies can be as general or specific as needed. The policy rules are compared against the traffic in sequence, so more specific rules must precede the more general ones.



Note: Refer to the Palo Alto Networks tech note, Controlling SSL Decryption, for instructions on managing SSL certificates to avoid certificate mismatch errors, and SSL Decryption Certificates for guidelines on how to develop policies to handle non-standard SSL implementations.

SSL forward proxy decryption requires the configuration of a trusted certificate that will be presented to the user if the server to which the user is connecting possesses a certificate signed by a CA trusted by the firewall. To configure this certificate, create a certificate on the **Device > Certificate Management > Certificates** page and then click the name of the certificate and check the **Forward Trust Certificate** check box. Refer to "Importing, Exporting and Generating Security Certificates" on page 86.



Note: Certain applications will not function if they are decrypted by the firewall. To prevent this from occurring, PAN-OS will not decrypt the SSL traffic for these applications and the decryption rule settings will not apply. For a list of these applications, refer to support article located at https:// live.paloaltonetworks.com/docs/DOC-1423.

Field	Description
General Tab	
Name	Enter a name to identify the rule (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required.
Description	Enter a description for the rule (up to 255 characters).
Tag	If you need to tag the policy, click Add to specify the tag.
	A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location.
Source Tab	
Source Zone	Click Add to choose source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.
	Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases.
Source Address	Click Add to add source addresses, address groups, or regions (default is any). Select from the drop-down list, or click the Address , Address Group , or Regions link at the bottom of the drop-down list, and specify the settings. Select the Negate check box to choose any address except the configured ones.

Table 83. Decryption Rule Settings

Field	Description
Source User	Click Add to choose the source users or groups of users subject to the policy. If you choose the pre-logon option, this will include remote users that are connected to the network using GlobalProtect, but are not logged into their system. When the Pre-logon option is configured on the Portal for GlobalProtect clients, any user who is not currently logged into their machine will be identified with the username pre-logon . You can then create policies for pre-logon users and the machines can be accessed as if they were logged in and connected to the domain.
Destination Tab	
Destination Zone	Click Add to choose destination zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151. Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing Sales and Public
	Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases.
Destination Address	Click Add to add destination addresses, address groups, or regions (default is any). Select from the drop-down list, or click the Address , Address Group , or Regions link at the bottom of the drop-down list, and specify the settings. Select the Negate check box to choose any address except the configured ones.
URL Category Tab	Select URL categories for the decryption rule.
	 Choose any to match any sessions regardless of the URL category.
	• To specify a category, click Add and select a specific category (including a custom category) from the drop-down list. You can add multiple categories. Refer to "Custom URL Categories" on page 242 for information on defining custom categories.
Options Tab	
Action	Select decrypt or no-decrypt for the traffic.
Туре	Select the type of traffic to decrypt from the drop-down list:
	• SSL Forward Proxy —Specifies that the policy will decrypt client traffic destined for an external server.
	• SSH Proxy —Specifies that the policy will decrypt SSH traffic. This option allows you to control SSH tunneling in policies by specifying the ssh-tunnel App-ID.
	• SSL Inbound Inspection —Specifies that the policy will decrypt SSL inbound inspection traffic.
Decryption Profile	Select an existing decryption profile, or create a new decryption profile. Refer to "Decryption Profiles" on page 248.

Table 83. Decryption Rule Settings (Continued)

Application Override Policies

To change how the firewall classifies network traffic into applications, you can specify application override policies. For example, if you want to control one of your custom applications, an application override policy can be used to identify traffic for that application according to zone, source and destination address, port, and protocol. If you have network applications that are classified as "unknown," you can create new application definitions for them (refer to "Defining Applications" on page 233).

Like security policies, application override policies can be as general or specific as needed. The policy rules are compared against the traffic in sequence, so the more specific rules must precede the more general ones.

Custom Application Definition with Application Override

Because the App-ID engine in PAN-OS classifies traffic by identifying the application-specific content in network traffic, the custom application definition cannot simply use a port number to identify an application. The application definition must also include traffic (restricted by source zone, source IP address, destination zone, and destination IP address).

To create a custom application with application override:

- 1. Define the custom application. Refer to "Defining Applications" on page 233. It is not required to specify signatures for the application if the application is used only for application override rules.
- 2. Define an application override policy that specifies when the custom application should be invoked. A policy typically includes the IP address of the server running the custom application and a restricted set of source IP addresses or a source zone.

Defining Application Override Policies

▶ Policies > Application Override

After creating a new rule, configure the rule by clicking the current field values and specifying the appropriate information, as described in the following table.

Field	Description
General Tab	
Name	Enter a name to identify the rule (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required.
Description	Enter a description for the rule (up to 255 characters).
Tag	If you need to tag the policy, click Add to specify the tag. A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location.

Table 84. Application Override Rule Settings

Field	Description
Source Tab	
Source Zone	Click Add to choose source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.
	Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases.
Source Address	Click Add to add source addresses, address groups, or regions (default is any). Select from the drop-down list, or click the Address , Address Group , or Regions link at the bottom of the drop-down list, and specify the settings. Select the Negate check box to choose any address except the configured ones.
Destination Tab	
Destination Zone	Click Add to choose destination zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.
	Multiple zones can be used to simplify management. For example, if you have three different internal zones (Marketing, Sales, and Public Relations) that are all directed to the untrusted destination zone, you can create one rule that covers all cases.
Destination Address	Click Add to add destination addresses, address groups, or regions (default is any). Select from the drop-down list, or click the Address , Address Group , or Regions link at the bottom of the drop-down list, and specify the settings. Select the Negate check box to choose any address except the configured ones.
Protocol/Application Tab	
Protocol	Select the protocol for which the application can be overridden.
Port	Enter the port number (0 to 65535) or range of port numbers (port1-port2) for the specified destination addresses. Multiple ports or ranges must be separated by commas.
Application	Select the override application for traffic flows that match the above rule criteria. When overriding to a custom application, there is no threat inspection that is performed. The exception to this is when you override to a pre-defined application that supports threat inspection.
	To define new applications, refer to "Defining Applications" on page 233).

Table 84. Application Override Rule Settings (Continued)

Captive Portal Policies

You can set up and customize a captive portal to direct user authentication by way of an authentication profile, an authentication sequence, or a certificate profile. Captive portal is used in conjunction with the User-ID Agent to extend user identification functions beyond the Active Directory domain. Users are directed to the portal and authenticated, thereby creating a user-to-IP address mapping.

Defining Captive Portal Policies

Policies > Captive Portal

Before you define captive portal policies, enable captive portal and configure captive portal settings on the **User Identification** page, as described in "Configuring the Firewall for User Identification" on page 286.

After doing so, configure capture portal policies by specifying the following information.

Field	Description
Name	Enter a name to identify the rule (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required.
Description	Enter a description for the rule (up to 255 characters).
Tag	If you need to tag the policy, click Add to specify the tag. A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location.
Source Tab	
Source	Specify the following information:
	• Choose a source zone if the policy needs to be applied to traffic coming from all interfaces in a given zone. Click Add to specify multiple interfaces or zones.
	• Specify the Source Address setting to apply the captive portal policy for traffic coming from specific source addresses. Select the Negate check box to choose any address except the configured ones. Click Add to specify multiple interfaces or zones.
Destination Tab	
Destination	Specify the following information:
	• Choose a destination zone if the policy needs to be applied to traffic to all interfaces in a given zone. Click Add to specify multiple interfaces or zones.
	• Specify the Destination Address setting to apply the captive portal policy for traffic to specific destination addresses. Select the Negate check box to choose any address except the configured ones. Click Add to specify multiple interfaces or zones.

Table 85. Captive Portal Rule Settings

Field	Description
Service/ URL Category Tab	
Service	Select services to limit to specific TCP and/or UDP port numbers. Choose one of the following from the drop-down list:
	 any—The selected services are allowed or denied on any protocol or port.
	• default —The selected services are allowed or denied only on the default ports defined by Palo Alto Networks. This option is recommended for allow policies.
	• Select—Click Add. Choose an existing service or choose Service or Service Group to specify a new entry. Refer to "Services" on page 239 and "Service Groups" on page 240.
URL Category	Select URL categories for the captive portal rule.
	 Choose any to apply the actions specified on the Service/Action tab regardless of the URL category.
	• To specify a category, click Add and select a specific category (including a custom category) from the drop-down list. You can add multiple categories. Refer to "Custom URL Categories" on page 242 for information on defining custom categories.
Service/Action Tab	
Action Setting	Choose an action to take:
	 web-form—Present a captive portal page for the user to explicitly enter authentication credentials.
	 no-captive-portal—Allow traffic to pass without presenting a captive portal page for authentication.
	• browser-challenge—Open an NT LAN Manager (NTLM) authentica- tion request to the user's web browser. The web browser will respond using the user's current login credentials.

 Table 85.
 Captive Portal Rule Settings (Continued)

DoS Protection Policies

DoS protection policies allow you to control the number of sessions between interfaces, zones, addresses, and countries based on aggregate sessions or source and/or destination IP addresses.

Defining DoS Protection Policies

Policies > DoS Protection

Use this page to add, edit, or delete DoS protection policy rules. To add a policy rule, click **Add** and then complete the following fields:

Field	Description
General Tab	
Name	Enter a name to identify the rule (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores. Only the name is required.
Description	Enter a description for the rule (up to 255 characters).
Tag	If you need to tag the policy, click Add to specify the tag. A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location.
Source Tab	
Source	 Specify the following information: Choose Interface from the Type drop-down list to apply the DoS policy to traffic coming from an interface or a group of interfaces. Choose Zone if the DoS policy needs to be applied to traffic coming from all interfaces in a given zone. Click Add to specify multiple interfaces or zones. Specify the Source Address setting to apply the DoS policy for traffic coming from specific source addresses. Select the Negate check box to choose any address except the configured ones. Click Add to specify multiple addresses. Specify the Source User setting to apply the DoS policy for traffic from specific users. Click Add to specify multiple users.

Table 86. DoS Rule Settings

Field	Description
Destination Tab	
Destination	Specify the following information:
	• Choose Interface from the Type drop-down list to apply the DoS policy to traffic coming from an interface or a group of interfaces. Choose Zone if the DoS policy needs to be applied to traffic coming from all interfaces in a given zone. Click Add to specify multiple interfaces or zones.
	• Specify the Destination Address setting to apply the DoS policy for traffic to specific destination addresses. Select the Negate check box to choose any address except the configured ones. Click Add to specify multiple addresses.
Option/Protection Tab	
Service	Select from the drop-down list to apply the DoS policy to only the configured services.
Action	Choose the action from the drop-down list:
	• Deny —Drop all traffic.
	• Allow—Permit all traffic.
	• Protect —Enforce protections supplied in the thresholds that are config- ured as part of the DoS profile applied to this rule.
Schedule	Select a pre-configured schedule from the drop-down list to apply the DoS rule to a specific date/time.
Log Forwarding	If you want to trigger forwarding of threat log entries to an external service—such as a syslog server or Panorama—select a log forwarding profile from the drop-down or click Profile to create a new one. Note that only traffic that matches an action in the rule will be logged and forwarded.
Aggregate	Select a DoS protection profile from the drop-down list to determine the rate at which you want to take action in response to DoS threats. The aggregate setting applies to the total of all traffic from the specified source to specified destination.
Classified	Select the check box and specify the following:
	• Profile —Select the profile from the drop-down list.
	• Address—Select whether to apply the rule to the source, destination, or source and destination IP addresses.
	If a classified profile is specified, the profile limitations are applied to a source IP address, destination IP address, or source and destination IP address pair. For example, you could specify a classified profile with a session limit of 100 and specify an Address setting of "source" in the rule. The result would be a limit of 100 sessions at any given time for that particular source IP address.

Table 60. Dos Kule senings (Commued	Table 86.	DoS Rule	Settings	(Continued
-------------------------------------	-----------	----------	----------	------------

Security Profiles

Each security policy can include specification of one or more security profiles, which provide additional protection and control.

You can also add threat exceptions to Anti-spyware and Vulnerability profiles. To make management of threat exceptions easier, you can add threat exceptions directly from the Monitor > Logs > Threat list. To add a threat exception to a profile, locate the threat in the Threat log and then click on the threat name. The Threat Details dialogue will appear, select one or more profiles in the left pane and click OK. This will add a threat exception for that signature in the selected threat profiles with the action of allow. To add a threat exception for specific IP addresses, add the IP addresses in the right pane and click OK. This will add a threat exception.

The following profile types are available:

- Antivirus profiles to protect against worms and viruses or block spyware downloads. Refer to "Antivirus Profiles" in the next section.
- Anti-spyware profiles to block attempts by spyware to access the network. Refer to "Anti-spyware Profiles" on page 213.
- Vulnerability protection profiles to stop attempts to exploit system flaws or gain unauthorized access to systems. Refer to "Vulnerability Protection Profiles" on page 215.
- URL filtering profiles to restrict access to specific web sites and web site categories. Refer to "URL Filtering Profiles" on page 217.
- File blocking profiles to block selected file types. Refer to "File Blocking Profiles" on page 220.
- Data filtering profiles that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall. Refer to "Data Filtering Profiles" on page 223.

In additional to individual profiles, you can create profile groups to combine profiles that are often applied together.



Note: You cannot delete a profile that is used in a security policy.

You can choose from the following actions when defining antivirus and anti-spyware profiles.

- **Default**—Takes the default action that is specified internally in the signature for each threat.
- Alert—Generates an alert for each application traffic flow. The alert is saved in the threat log.
- **Block**—Drops the application traffic.
- Allow—Permits the application traffic.

The following actions are available when defining custom Spyware and Vulnerability objects:

• Alert—Generates an alert for each application traffic flow. The alert is saved in the threat log.

- Drop Packets—Keeps all packets from continuing past the firewall.
- **Reset Both**—Resets the client and server.
- **Reset Client**—Resets the client.
- **Reset Server**—Resets the server.
- **Block-IP**—This action blocks traffic from either a source or a source-destination pair (configurable) for a specified period of time.

Antivirus Profiles

Objects > Security Profiles > Antivirus

A security policy can include specification of an antivirus profile to identify which applications are inspected for viruses and the action taken when a virus is detected. The profile can also specify actions to block spyware downloads. The default profile inspects all of the listed protocol decoders for viruses, generates alerts for Simple Mail Transport Protocol (SMTP), Internet Message Access Protocol (IMAP), and Post Office Protocol Version 3 (POP3), and takes the default action for other applications (alert or deny), depending on the type of virus detected.

Customized profiles can be used to minimize antivirus inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms. For information on action types, refer to "Security Profiles" on page 211.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of antivirus profiles when defining security policies. The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, periods, and underscores.
Description	Enter a description for the profile (up to 255 characters).
Antivirus Tab	
Packet Capture	Select the check box if you want to capture identified packets.
Decoders and Actions	For each type of traffic that you want to inspect for viruses, select an action from the drop-down list. You can also take specific action based on signatures created by WildFire. For more information on WildFire, refer to "About WildFire" on page 421.

Table 87. Antivirus Profile Settings

Field	Description
Applications Exceptions and Actions	Identify applications that will be exceptions to the antivirus rule. For example, to block all HTTP traffic except for a specific application, you can define an antivirus profile for which the application is an exception. Block is the action for the HTTP decoder, and Allow is the exception for the application.
	To find an application, start typing the application name in the text box. A matching list of applications is displayed, and you can make a selection. The application is added to the table, and you can assign an action.
	For each application exception, select the action to be taken when the threat is detected.
Virus Exception Tab	
Threat ID	Use this tab if you want the system to ignore specific threats. Exceptions that are already specified are listed. You can add additional threats by entering the threat ID and clicking Add . Threat IDs are presented as part of the threat log information. Refer to "Viewing the Logs" on page 264.

Table 87. Antivirus Profile Settings (Continued)

Anti-spyware Profiles

Objects > Security Profiles > Anti-spyware

A security policy can include specification of an anti-spyware profile for "phone home" detection (detection of traffic from installed spyware). The default anti-spyware profile detects phone-home protection for all severity levels except the low and informational levels.

Customized profiles can be used to minimize anti-spyware inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms.

The Exceptions settings allows you to change the response to a specific signature. For example, you can block all packets that match a signature, except for the selected one, which generates an alert.

The **DNS Signatures** settings provides an additional method of identifying infected hosts on a network. These signatures detect specific DNS lookups for host names that have been associated with malware. The DNS signatures can be configured to allow, alert, or (default) block when these queries are observed, just as with regular antivirus signatures. Additionally, hosts that perform DNS queries for malware domains will appear in the botnet report. DNS signatures are downloaded as part of the antivirus updates.

The **Anti-spyware** page presents a default set of columns. Additional columns of information are available by using the column chooser. Click the arrow to the right of a column header and select the columns from the Columns sub-menu. For more information, refer to "Using Tables on Configuration Pages" on page 26.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of anti-spyware profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, periods, and underscores.
Description	Enter a description for the profile (up to 255 characters).
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.
Rules Tab	
Rule Name	Specify the rule name.
Threat Name	Enter any to match all signatures, or enter text to match any signature containing the entered text as part of the signature name.
Severity	Choose a severity level (critical, high, medium, low, or informational).
Action	Choose an action (Default , Alert , Allow , or Drop) for each threat.
Packet Capture	Select the check box if you want to capture identified packets.
Exceptions Tab	
Exceptions	Select the Enable check box for each threat for which you want to assign an action, or select All to respond to all listed threats. The list depends on the selected host, category, and severity. If the list is empty, there are no threats for the current selections.
	Use the IP Address Exemptions column to add IP address filters to a threat exception. If IP addresses are added to a threat exception, the threat exception action for that signature will only be taken over the rule's action if the signature is triggered by a session having either the source or destination IP matching an IP in the exception. You can add up to 100 IP addresses per signature. With this option, you do not have to create a new policy rule and new vulnerability profile to create an exception for a specific IP address.
DNS Signature Tab	
Infected host DNS Signature	Choose an action to be taken when DNS lookups are made to known malware sites (Alert, Allow, or default (Block)).
Packet Capture	Select the check box if you want to capture identified packets.
Threat ID	Manually enter DNS signature exceptions (range 4000000-4999999).

Table 88. Anti-spyware Profile Settings

Vulnerability Protection Profiles

Objects > Security Profiles > Vulnerability Protection

A security policy can include specification of a vulnerability protection profile that determines the level of protection against buffer overflows, illegal code execution, and other attempts to exploit system vulnerabilities. The default profile protects clients and servers from all known critical, high, and medium-severity threats.

Customized profiles can be used to minimize vulnerability checking for traffic between trusted security zones, and to maximize protection for traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms. To apply vulnerability protection profiles to security policies, refer to "Security Policies" on page 187.

The Rules settings specify collections of signatures to enable, as well as actions to be taken when a signature within a collection is triggered.

The Exceptions settings allows you to change the response to a specific signature. For example, you can block all packets that match a signature, except for the selected one, which generates an alert. The **Exception** tab supports filtering functions.

The **Vulnerability Protection** page presents a default set of columns. Additional columns of information are available by using the column chooser. Click the arrow to the right of a column header and select the columns from the Columns sub-menu. For more information, refer to "Using Tables on Configuration Pages" on page 26.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of vulnerability protection profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, periods, and underscores.
Description	Enter a description for the profile (up to 255 characters).
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.
Rules Tab	
Rule Name	Specify a name to identify the rule.
Threat Name	Specify a text string to match. The firewall applies a collection of signatures to the rule by searching signature names for this text string.
Action	Choose the action (Alert, Allow, Default, or Block) to take when the rule is triggered. The Default action is based on the pre-defined action that is part of each signature provided by Palo Alto Networks. To view the default action for a signature, navigate to Objects > Security Profiles > Vulnerability Protection and click Add or select an existing profile. Click the Exceptions tab and then click Show all signatures. A list of all signatures will displayed and you will see an Action column.
Host	Specify whether to limit the signatures for the rule to those that are client side, server side, or either (any).
Packet Capture	Select the check box if you want to capture the packet that triggered the rule.

Table 89. Vulnerability Protection Profile Settings

Field	Description
Category	Select a vulnerability category if you want to limit the signatures to those that match that category.
CVE List	Specify common vulnerabilities and exposures (CVEs) if you want to limit the signatures to those that also match the specified CVEs.
	Each CVE is in the format CVE- <i>yyyy-xxxx</i> , where <i>yyyy</i> is the year and <i>xxxx</i> is the unique identifier. You can perform a string match on this field. For example, to find vulnerabilities for the year 2011, enter "2011".
Vendor ID	Specify vendor IDs if you want to limit the signatures to those that also match the specified vendor IDs.
	For example, the Microsoft vendor IDs are in the form MSyy-xxx, where <i>yy</i> is the two-digit year and xxx is the unique identifier. For example, to match Microsoft for the year 2009, enter "MS09".
Severity	Select severities to match (informational , low , medium , high , or critical) if you want to limit the signatures to those that also match the specified severities.
Exceptions Tab	
Threats	Select the Enable check box for each threat for which you want to assign an action, or select All to respond to all listed threats. The list depends on the selected host, category, and severity. If the list is empty, there are no threats for the current selections.
	Choose an action from the drop-down list box, or choose from the Action drop-down at the top of the list to apply the same action to all threats. If the Show All check box is selected, all signatures are listed. If the Show All check box is not selected, only the signatures that are exceptions are listed.
	Select the Packet Capture check box if you want to capture identified packets.
	The vulnerability signature database contains signatures that indicate a brute force attack; for example, Threat ID 40001 triggers on an FTP brute force attack. Brute-force signatures trigger when a condition occurs in a certain time threshold. The thresholds are pre-configured for brute force
	signatures, and can be changed by clicking the pencil icon Z next to the threat name on the Vulnerability tab (with the Custom option selected). You can specify the number of hits per unit of time and whether the threshold applies to source, destination, or source-and-destination.
	Thresholds can be applied on a source IP, destination IP or a combination of source IP and destination IP.
	Note: The default action is shown in parentheses. The CVE column shows identifiers for common vulnerabilities and exposures (CVE). These unique, common identifiers are for publicly known information security vulnerabilities.
	Use the IP Address Exemptions column to add IP address filters to a threat exception. If IP addresses are added to a threat exception, the threat exception action for that signature will only be taken over the rule's action if the signature is triggered by a session having either the source or destination IP matching an IP in the exception. You can add up to 100 IP addresses per signature. With this option, you do not have to create a new policy rule and new vulnerability profile to create an exception for a specific IP address.

 Table 89.
 Vulnerability Protection Profile Settings (Continued)
URL Filtering Profiles

Objects > Security Profiles > URL Filtering

A security policy can include specification of a URL filtering profile that blocks access to specific web sites and web site categories, or generates an alert when the specified web sites are accessed (a URL filtering license is required). You can also define a "block list" of web sites that are always blocked (or generate alerts) and an "allow list" of web sites that are always allowed. The web categories are predefined by Palo Alto Networks.

To apply URL filtering profiles to security policies, refer to "Security Policies" on page 187. To create custom URL categories with your own lists of URLs, refer to "Custom URL Categories" on page 242.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of URL filtering profiles when defining security policies. The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Description	Enter a description for the profile (up to 255 characters).
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.
Action on License	Select the action to take if the URL filtering license expires:
Expiration	• Block —Blocks access to all web sites in the block list or the selected categories.
	• Allow—Allows access to all web sites.
Dynamic URL Filtering	Select to enable dynamic URL categorization.
	URL categorization takes advantage of a URL filtering database on the firewall that lists the most popular URLs and other URLs for malicious categories. The URL filtering database may be able to resolve requests that the local database is unable to categorize. The default is enabled when using the BrightCloud database. When using the PAN-DB, this option is enabled by default and is not configurable.
	To configure the system response when a URL remains unresolved after a 5 second timeout period, use the Category and Action settings in this window (see Category Action later in this table). Select the action for the category "Not resolved URL."
Log Container Page Only	Select the check box to log only the URLs that match the content type that is specified. The default is enabled.

Table 90. URL Filtering Profile Settings

Field	Description
Block List	Enter the IP addresses or URL path names of the web sites that you wa to block or generate alerts on. Enter each URL one per line.
	IMPORTANT: You must omit the "http and https" portion of the URI when adding web sites to the list.
	Entries in the block list are an exact match and are case-insensitive. For example, "www.paloaltonetworks.com" is different from "paloaltonetworks.com". If you want to block the entire domain, you should include both "*.paloaltonetworks.com" and "paloaltonetworks.com".
	Examples:
	• www.paloaltonetworks.com
	• 198.133.219.25/en/US
	Block and allow lists support wildcard patterns. The following character are considered separators:
	<pre> ' ' ' ' ' ' ' ' ' ' ' ' ' ' ' ' ' ' '</pre>
	www.y*.com
Action	Select the action to take when a web site in the block list is accessed.
	 alert—Allow the user to access the web site, but add an alert to the Ullog.
	• block —Block access to the web site.
	 continue—Allow the user to access the blocked page by clicking Continue on the block page.
	 override—Allow the user to access the blocked page after entering a password. The password and other override settings are specified in t URL Admin Override area of the Settings page. Refer to Table 1 in tl "Defining Management Settings" on page 30.

Table 90. URL Filtering Profile Settings (Continued)

Field	Description
Allow List	Enter the IP addresses or URL path names of the web sites that you want to allow or generate alerts on. Enter each URL one per line.
	IMPORTANT: You must omit the "http and https" portion of the URLs when adding web sites to the list.
	Entries in the allow list are an exact match and are case-insensitive. For example, "www.paloaltonetworks.com" is different from "paloaltonetworks.com". If you want to allow the entire domain, you should include both "*.paloaltonetworks.com" and "paloaltonetworks.com".
	Examples:
	www.paloaltonetworks.com
	• 198.133.219.25/en/US
	Block and allow lists support wildcard patterns. The following characters are considered separators:
	/ ? & = ;
	+ Every substring that is separated by the characters listed above is considered a token. A token can be any number of ASCII characters that does not contain any separator character or *. For example, the following patterns are valid:
	.yahoo.com (Tokens are: "", "yahoo" and "com") www.*.com (Tokens are: "www", "*" and "com")
	<pre>www.yahoo.com/search=* (Tokens are: "www", "yahoo", "com", "search", "*") The following patterns are invalid because the character "*" is not the only character in the token</pre>
	าน (1) (1) (1) (1) (1) (1) (1) (1) (1) (1)
	This list takes precedence over the selected web site categories.
Category/Action	For each category, select the action to take when a web site of that category is accessed.
	• alert —Allow the user to access the web site, but add an alert to the URL log.
	• allow —Allow the user to access the web site.
	• block —Block access to the web site.
	 continue—Allow the user to access the blocked page by clicking Con- tinue on the block page.
	• override —Allow the user to access the blocked page after entering a password. The password and other override settings are specified in the URL Admin Override area of the Settings page. Refer to Table 1 in the "Defining Management Settings" on page 30.
Check URL Category	Click to access the web site where you can enter a URL or IP address to view categorization information.

Table 90. URL Filtering Profile Settings (Continued)

File Blocking Profiles

Objects > Security Profiles > File Blocking

A security policy can include specification of a file blocking profile that blocks selected file types from being uploaded and/or downloaded, or generates an alert when the specified file types are detected. Table 92 lists the supported file formats at the time of this publication. However, because new file type support can be added in a content update, for the most up-to-date list, click **Add** in the **File Types** field of the File Blocking Profile dialog.

To apply file blocking profiles to security policies, refer to "Security Policies" on page 187.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of file blocking profiles when defining security policies. The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Description	Enter a description for the profile (up to 255 characters).
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.
Rules	Define one or more rules to specify the action taken (if any) for the selected file types. To add a rule, specify the following and click Add :
	• Name—Enter a rule name (up to 31 characters).
	• Applications—Select the applications the rule applies to or select any.
	• File Types—Select the file types for which you want to block or generate alerts.
	• Direction —Select the direction of the file transfer (Upload, Download, or Both).
	 Action—Select the action taken when the selected file types are detected:
	 alert—An entry is added to the threat log.
	– block —The file is blocked.
	 continue—A message to the user indicates that a download has been requested and asks the user to confirm whether to continue. The purpose is to warn the user of a possible unknown download (also known as a drive-by-download) and to give the user the option of continuing or stopping the download.
	Note: When you create a file blocking profile with the action continue or continue-and-forward (used for WildFire forwarding), you can only choose the application web-browsing . If you choose any other application, traffic that matches the security policy will not flow through the firewall due to the fact that the users will not be prompted with a continue page.
	 forward—The file is automatically sent to WildFire. continue-and-forward—A continue page is presented, and the file is sent to WildFire (combines the continue and forward actions).

 Table 91. File Blocking Profile Settings

Field	Description
apk	Android application package file
avi	Video file based on Microsoft AVI (RIFF) file format
avi-divx	AVI video file encoded with the DivX codec
avi-xvid	AVI video file encoded with the XviD codec
bat	MS DOS Batch file
bmp-upload	Bitmap image file (upload only)
cab	Microsoft Windows Cabinet archive file
cdr	Corel Draw file
class	Java bytecode file
cmd	Microsoft command file
dll	Microsoft Windows Dynamic Link Library
doc	Microsoft Office Document
docx	Microsoft Office 2007 Document
dpx	Digital Picture Exchange file
dsn	Database Source Name file
dwf	Autodesk Design Web Format file
dwg	Autodesk AutoCAD file
edif	Electronic Design Interchange Format file
encrypted-doc	Encrypted Microsoft Office Document
encrypted-docx	Encrypted Microsoft Office 2007 Document
encrypted-office2007	Encrypted Microsoft Office 2007 Fil
encrypted-pdf	Encrypted Adobe PDF Document
encrypted-ppt	Encrypted Microsoft Office PowerPoint
encrypted-pptx	Encrypted Microsoft Office 2007 PowerPoint
encrypted-rar	Encrypted rar file
encrypted-xls	Encrypted Microsoft Office Excel
encrypted-xlsx	Encrypted Microsoft Office 2007 Excel
encrypted-zip	Encrypted zip file
exe	Microsoft Windows Executable
flv	Adobe Flash Video file
gds	Graphics Data System file
gif-upload	GIF image file (upload only)
gzip	Files compressed with gzip utility
hta	HTML Application file
iso	Disc Image file based on ISO-9660 standard

Table 92. Supported File Formats for File Blocking

Field	Description
iwork-keynote	Apple iWork Keynote documents
iwork-numbers	Apple iWork Numbers documents
iwork-pages	Apple iWork Pages documents
jar	Java ARchive
jpeg-upload	JPG/JPEG image file (upload only)
lnk	Microsoft Windows file shortcut
lzh	File compressed with lha/lzh utility/algorithm
mdb	Microsoft Access Database file
mdi	Microsoft Document Imaging file
mkv	Matroska Video file
mov	Apple Quicktime Movie file
mp3	MP3 audio file
mp4	MP4 audio file
mpeg	Movie file using MPEG-1 or MPEG-2 compression
msi	Microsoft Windows Installer package file
msoffice	Microsoft Office File (doc, xls, ppt, pub, pst)
ocx	Microsoft ActiveX file
pdf	Adobe Portable Document file
PE	Microsoft Windows Portable Executable (exe, dll, com, scr, ocx, cpl, sys, drv, tlb)
pgp	Security key or digital signature encrypted with PGP software
pif	Windows Program Information File containing executable instructions
pl	Perl Script file
png-upload	PNG image file (upload only)
ppt	Microsoft Office PowerPoint Presentation
pptx	Microsoft Office 2007 PowerPoint Presentation
psd	Adobe Photoshop Document
rar	Compressed file created with winrar
reg	Windows Registry file
rm	RealNetworks Real Media file
rtf	Windows Rich Text Format document file
sh	Unix Shell Script file
stp	STandard for the Exchange of Product model data 3D graphic file
tar	Unix tar archive file
tdb	Tanner Database (www.tannereda.com)
tif	Windows Tagged Image file

 Table 92.
 Supported File Formats for File Blocking (Continued)

Field	Description
torrent	BitTorrent file
wmf	Windows Metafile to store vector images
wmv	Windows Media Video file
wri	Windows Write document file
wsf	Windows Script file
xls	Microsoft Office Excel
xlsx	Microsoft Office 2007 Excel
zcompressed	Compressed Z file in Unix, decompressed with uncompress
zip	Winzip/pkzip file

Table 92. Supported File Formats for File Blocking (Continued)

Data Filtering Profiles

Objects > Security Profiles > Data Filtering

A security policy can include specification of a data filtering profile to help identify sensitive information such as credit card or social security numbers and prevent the sensitive information from leaving the area protected by the firewall.

To apply data filtering profiles to security policies, refer to "Security Policies" on page 187.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of log forwarding profiles when defining security policies. The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Description	Enter a description for the profile (up to 255 characters).
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.
Data Capture	Select the check box to automatically collect the data that is blocked by the filter.

Table 93. Data Filtering Profile Settings



Note: Specify a password for Manage Data Protection on the *Settings* page to view your captured data. Refer to "Defining Management Settings" on page 30.

To add a data pattern, click **Add** and specify the following information.

Field	Description
Data Pattern	Choose an existing data pattern from the Data Pattern drop-down list, or configure a new pattern by choosing Data Pattern from the list and specifying the following information:
	• Name—Configure a name for the data pattern.
	• Description —Configure a description for the data pattern.
	 Shared—Select this option if the data pattern object will be shared across multiple virtual systems.
	• Weight—Specify unit values for the specified patterns to use in calculating thresholds. For instance, if you designate a weight of 5 for SSN#, every instance of a SSN pattern will increment the threshold by 5. In other words, the detection of ten SSN patterns will result in 10 x 5 (weight) = 50.
	- CC# —Specify a weight for the credit card field (range 0-255).
	 SSN#—Specify a weight for the social security number field, where the field includes dashes, such as 123-45-6789 (range 0-255, 255 is highest weight).
	 SSN# (without dash)—Specify a weight for the social security number field, where the entry is made without dashes, such as 123456789 (range 0-255, 255 is highest weight).
	• Custom Patterns —To match a custom data pattern for the traffic that is subject to this profile, create a custom data pattern by clicking Add and specifying the pattern name, regular expression (regex) to match, and weight (0-255, 255 is highest weight). You can add multiple match expressions to the same data pattern profile.
Applications	Specify the applications to include in the filtering rule:
	• Choose any to apply the filter to all of the listed applications. This selection does not block all possible applications, just the listed ones.
	 Click Add to specify individual applications.
File Types	Specify the file types to include in the filtering rule:
91	• Choose any to apply the filter to all of the listed file types. This selection does not block all possible file types, just the listed ones.
	 Click Add to specify individual file types.
Direction	Specify whether to apply the filter in the upload direction, download direction, or both.
Alert Threshold	Specify the value that will trigger an alert. For example, if you have a threshold of 100 with a SSN weight of 5, the rule will need to detect at least 20 SSN patterns before the rule will be triggered (20 instances x 5 weight = 100).
Block Threshold	Specify the value that will trigger a block. For example, if you have a threshold of 100 with a SSN weight of 5, the rule will need to detect at least 20 SSN patterns before the rule will be triggered (20 instances x 5 weight = 100).

Table 94. Data Pattern Settings

DoS Profiles

▶ Objects > Security Profiles > DoS Protection

A DoS protection policy can include specification of a DoS profile to protect against DoS attacks and take protective action in response to rule matches. The DoS profile specifies the types of actions and the matching criteria.

To apply DoS profiles to DoS policies, refer to "DoS Protection Policies" on page 209.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of log forwarding profiles when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.
Description	Enter a description of the profile (up to 255 characters).
Туре	Specify one of the following profile types:
	• aggregate —Apply the DoS thresholds configured in the profile to all packets that match the rule criteria on which this profile is applied. For example, an aggregate rule with a SYN flood threshold of 10000 packets per second (pps) counts all packets that hit that particular DoS rule.
	• classified —Apply the DoS thresholds configured in the profile to all packets satisfying the classification criterion (source IP, destination IP or source-and-destination IP).
Flood Protection Tab	
Syn Flood subtab UDP Flood subtab	Select the check box to enable SYN flood protection, and specify the following settings:
ICMP Flood subtab	• Choice—(SYN Flood only) Choose from the following options:
Other subtab	 Random early drop—Drop packets randomly before the overall DoS limit is reached.
	 SYN cookies—Use SYN cookies to generate acknowledgments so that it is not necessary to drop connections in the presence of a SYN flood attack.
	• Alarm Rate—Specify the rate (pps) at which a DoS alarm is generated (range 0-2000000 pps, default 10000 pps).
	• Activate Rate—Specify the rate (pps) at which a DoS response is activated (range 0-2000000 pps, default 10000 pps).
	• Maximal Rate—Specify the rate at which packets will be dropped or blocked.
	• Block Duration —Specify the length of time (seconds) during which the offending packets will be denied. Packets arriving during the block duration do not count towards triggered alerts.
	Note: When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session.

Table 95. DoS Profile Settings

Field	Description
Resources Protection Tab	
Sessions	Select the check box to enable resources protection.
Max Concurrent Limit	Specify the maximum number of concurrent sessions. If the DoS profile type is aggregate, this limit applies to the entire traffic hitting the DoS rule on which the DoS profile is applied. If the DoS profile type is classified, this limit applies to the entire traffic on a classified basis (source IP, destination IP or source-and-destination IP) hitting the DoS rule on which the DoS profile is applied.

Table 95. DoS Profile Settings (Continued)

Other Policy Objects

Policy objects are the elements that enable you to construct, schedule, and search for policies. The following element types are supported:

- Addresses and address groups to determine the scope of the policy. Refer to "Addresses and Address Groups" in the next section.
- Applications and application groups that allow you to specify how software applications are treated in policies. Refer to "Applications and Application Groups" on page 231.
- Application filters that allow you to simplify searches. Refer to "Application Filters" on page 238.
- Services and service groups to limit the port numbers. Refer to "Services" on page 239.
- Data patterns to define categories of sensitive information for data filtering policies. Refer to "Data Patterns" on page 240.
- Custom URL categories that contain your own lists of URLs to include as a group in URL filtering profiles. Refer to "Custom URL Categories" on page 242.
- Spyware and vulnerability threats to allow for detailed threat responses. Refer to "Security Profile Groups" on page 246.
- Log forwarding to specify log settings. Refer to "Log Forwarding" on page 247.
- Schedules to specify when policies are active. Refer to "Schedules" on page 250.

Addresses and Address Groups

To define security policies for specific source or destination addresses, you must first define the addresses and address ranges. Addresses requiring the same security settings can be combined into address groups that you can refer to as a unit.

Defining Address Ranges

► Objects > Addresses

To define security policies for specific source or destination addresses, you must first define the addresses and address ranges. Addresses requiring the same security settings can be combined into address groups to simplify policy creation (refer to "Defining Address Groups" on page 230).

Field	Description
Name	Enter a name that describes the addresses to be defined (up to 63 characters). This name appears in the address list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow use by all virtual systems.
Description	Enter a description for the object (up to 255 characters).
Туре	Specify an IPv4 or IPv6 address or address range, FQDN, or Dynamic.
	IP Netmask:
	Enter the IPv4 or IPv6 address or IP address range using the following notation:
	<i>ip_address/mask</i> or <i>ip_address</i>
	where the <i>mask</i> is the number of significant binary digits used for the network portion of the address.
	Example:
	"192.168.80.150/32" indicates one address, and "192.168.80.0/24" indicates all addresses from 192.168.80.0 through 192.168.80.255.
	Example:
	"2001:db8:123:1::1" or "2001:db8:123:1::/64"
	IP Range:
	To specify an address range, select IP Range , and enter a range of addresses. The format is:
	ip_address-ip_address
	where each address can be IPv4 or IPv6.
	Example:
	"2001:db8:123:1::1 - 2001:db8:123:1::22"

Table 90. New Address Setting	Table 9	96.	New	Address	Setting
-------------------------------	---------	-----	-----	---------	---------

Field	Description				
Type (continued)	FQDN:				
	To specify an address using the FQDN, select FQDN and enter the domain name.				
	The FQDN initially resolves at commit time. Entries are subsequently refreshed when the firewall performs a check every 30 minutes; all changes in the IP address for the entries are picked up at the refresh cycle.				
	The FQDN is resolved a proxy is configured. Proxy" on page 173.	d by the system DNS server or a DNS proxy object. . For information about DNS proxy, refer to "DNS			
	Dynamic:				
	unique and is not lim and identifier specifie policies and the ident Dynamic" is used to r the XMP API script.	ited to a UUID standard. Once you have the name ed, you will use the "name" of the object in your tifier (the field located to the right of "Type: map the unique IP address to an address object in			
Name Test1					
	Name	: Test1			
	Name Description	Test1 Test address dynamic object.			
	Name Description Type				
	Name Description Type Iden	e Test 1 Test address dynamic object. Dynamic TesObjectID Enter an identifier for the dynamic object as it is used in the XML API. An identifier can be any string, but has to be unique.			

Table 96. New Address Settings (Continued)

This dynamic address object feature is useful in environments where host IP addresses change frequently; like in virtual environments where hosts are often added and removed and where failovers may occur requiring a change in IP address. By using a name for the dynamic address object for policy matching instead of static IP addresses or networks, the addresses can be dynamically updated without modifying the firewall configuration.

Each Dynamic Address Object can contain up to 256 IP addresses.

Defining Address Groups

► Objects > Address Groups

To simplify the creation of security policies, addresses requiring the same security settings can be combined into address groups.

Field	Description
Name	Enter a name that describes the address group (up to 63 characters). This name appears in the address list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow use by all virtual systems.
Addresses	Click Add and select addresses and/or other address groups to be included in this group.

Defining Regions

Objects > Regions

The firewall supports creation of policy rules that apply to specified countries or other regions. The region is available as an option when specifying source and destination for security policies, decryption policies, and DoS policies. You can choose from a standard list of countries or use the region settings described in this section to define custom regions to include as options for security policy rules.

Table 98.	New	Address	Settings
-----------	-----	---------	----------

Field	Description
Name	Enter a name that describes the region (up to 31 characters). This name appears in the address list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Geo Location	To specify latitude and longitude, select the check box and values (<i>xxx.xxxxx</i> format). This information is used in the traffic and threat maps for App-Scope. Refer to "Using App-Scope" on page 256.
Addresses	Specify an IP address, range of IP addresses, or subnet to identify the region, using any of the following formats:
	<i>x.x.x.x</i>
	x.x.x.x-y.y.y.y
	x.x.x.x/n

Applications and Application Groups

► Objects > Applications

The **Applications** page lists various attributes of each application definition, such as the application's relative security risk (1 to 5). The risk value is based on criteria such as whether the application can share files, is prone to misuse, or tries to evade firewalls. Higher values indicate higher risk.

The top application browser area of the page lists the attributes that you can use to filter the display. The number to the left of each entry represents the total number of applications with that attribute.

You can perform any of the following functions on this page:

• To apply application filters, click an item that you want to use as a basis for filtering. For example, to restrict the list to the Networking category, click **Networking** and the list will only show networking applications.

Search	q	L	Custom Only	ear Filters	1244 match
Category 🔺	Subcategory 📥		Technology 🔺	Risk 📥	Characteristic 🔺
241 business-systems	31 audio-streaming	^	476 browser-based	317 1	499 Evasive
333 collaboration	11 auth-service		452 dient-server	242 2	394 Excessive Bandwid
196 general-internet	15 database		197 network-protocol	294 3	264 Prone to Misuse
173 media	56 email		117 peer-to-peer	258 4	612 Transfers Files
299 networking	33 encrypted-tunnel			133 5	249 Tunnels Other App
2 unknown	19 erp-crm				266 Used by Malware
	154 file-sharing				733 Vulnerability
	46 gaming	~			779 Widely used

• To filter on additional columns, select an entry in the other columns. The filtering is successive: first category filters are applied, then subcategory filters, then technology filters, then risk, and finally characteristic filters.

For example, the next figure shows the result of applying a category, subcategory, and risk filter. In applying the first two filters, the **Technology** column is automatically restricted to the technologies that are consistent with the selected category and sub category, even though a technology filter has not been explicitly applied.

Each time a filter is applied, the list of applications in the lower part of the page is automatically updated, as shown in the following figure. Any saved filters can be viewed in **Objects > Application Filters**.

Search	٩	Custom Only	ar Filters	5 matching applications
Category 🔺	Subcategory 📥	Technology 🔺	Risk 📥	Characteristic 🔺
5 collaboration	56 email	5 browser-based	2 1	4 Transfers Files
	91 instant-messaging		1 2	1 Tunnels Other Apps
	30 internet-conferencing		2 3	4 Vulnerability
	5 social-business			2 Widely used
	64 social-networking			
	52 voip-video			
	35 web-posting			
Name	Category	Subcategory		Risk Technology
💷 blackboard	collaboration	social-business		t browser-based
III clearspace	collaboration	social-business		3 browser-based
III projectplace	collaboration	social-business		2 browser-based
💷 rypple	collaboration	social-business	l	browser-based
💷 sharepoint (1 out of 6 sł	nown)			
🦾 💀 sharepoint-base	collaboration	social-business		3 browser-based

• To search for a specific application, enter the application name or description in the **Search** field, and press **Enter**. The application is listed, and the filter columns are updated to show statistics for the applications that matched the search.

A search will match partial strings. When you define security policies, you can write rules that apply to all applications that match a saved filter. Such rules are dynamically updated when a new application is added through a content update that matches the filter.

• Click an application name to view additional details about the application, as described in the following table. You can also customize risk and timeout values, as described in the following table.

ltem	Description
Name	Name of the application.
Description	Description of the application (up to 255 characters).
Additional Information	Links to web sources (Wikipedia, Google, and Yahoo!) that contain additional information about the application.
Standard Ports	Ports that the application uses to communicate with the network.
Capable of File Transfer	Indication of whether the application is able to transfer files.
Used by Malware	Indication of whether the application is used by malware.
Excessive Bandwidth Use	Indication of whether the application uses too much bandwidth so that network performance may be compromise.
Evasive	Indication of whether the application attempts to evade firewalls.
Widely used	Indication of whether the effects of the application are wide- ranging.
Has Known Vulnerabilities	Indication of whether the application has any currently known vulnerabilities.
Tunnels Other Applications	Indication of whether the application can carry other applications within the messages that it sends.
Depends on Applications	List of other applications that are required for this application to run.
Category	Application category.
Subcategory	Application sub category.
Technology	Application technology.
	Assigned risk of the application.
Risk	To customize this setting, click the Customize link, enter a value (1-5), and click OK .
Prone to Misuse	Indication of whether the application tends to attract misuse.
Session Timeout	Period of time (seconds) required for the application to timeout due to inactivity (1-604800 seconds). This timeout is for protocols other than TCP or UDP. For TCP and UDP, refer to the next rows in this table.
	To customize this setting, click the Customize link, enter a value (seconds), and click OK .

Table 99. Application Details

ltem	Description
	Timeout for terminating a TCP application flow (1-604800 seconds).
TCP Timeout (seconds)	To customize this setting, click the Customize link, enter a value (seconds), and click OK .
	Timeout for terminating a UDP application flow (1-604800 seconds).
UDP Timeout (seconds):	To customize this setting, click the Customize link, enter a value (seconds), and click OK .

Table 99. Application Details (Continued)

When the firewall is not able to identify an application using the application ID, the traffic is classified as unknown: unknown-tcp or unknown-udp. This behavior applies to all unknown applications except those that fully emulate HTTP. For more information, refer to "Identifying Unknown Applications and Taking Action" on page 276.

You can create new definitions for unknown applications and then define security policies for the new application definitions. In addition, applications that require the same security settings can be combined into application groups to simplify the creation of security policies.

Defining Applications

► Objects > Applications

Use the **Applications** page to add new applications for the firewall to evaluate when applying policies.

Field	Description
Configuration Tab	
Name	Enter the application name (up to 31 characters). This name appears in the applications list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, periods, hyphens, and underscores. The first character must be a letter.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the application to be shared by all virtual systems.
Description	Enter a description of the application for general reference (up to 255 characters).
Category	Select the application category, such as email or database. For a description of each category, refer to "Application Categories and Subcategories" on page 437. The category is used to generate the Top Ten Application Categories chart and is available for filtering (refer to "Using the Application Command Center" on page 253).
Subcategory	Select the application subcategory, such as email or database. For a description of each sub category, refer to "Application Categories and Subcategories" on page 437. The sub category is used to generate the Top Ten Application Categories chart and is available for filtering (refer to "Using the Application Command Center" on page 253).
Technology	Select the technology for the application. For a description of each technology, refer to "Application Technologies" on page 439.

Table 100. New Application Settings

Field	Description
Parent App	Specify a parent application for this application. This setting applies when a session matches both the parent and the custom applications; however, the custom application is reported because it is more specific.
Risk	Select the risk level associated with this application (1=lowest to 5=highest).
Characteristics	Select the application characteristics that may place the application at risk. For a description of each characteristic, refer to "Application Characteristics" on page 439.
Advanced Tab	
Port	If the protocol used by the application is TCP and/or UDP, select Port and enter one or more combinations of the protocol and port number (one entry per line). The general format is:
	<pre><pre>cprotocol><<pre>cprotocol><<pre>cprotocol></pre><pre>where the <port> is a single port number, or dynamic for dynamic port assignment.</port></pre></pre></pre></pre>
	Examples: TCP/dynamic or UDP/32.
	This setting applies when using app-default in the Service column of a security rule.
IP Protocol	To specify an IP protocol other than TCP or UDP, select IP Protocol , and enter the protocol number (1 to 255).
ІСМР Туре	To specify an Internet Control Message Protocol version 4 (ICMP) type, select ICMP Type and enter the type number (range 0-255).
ІСМР6 Туре	To specify an Internet Control Message Protocol version 6 (ICMPv6) type, select ICMP6 Type and enter the type number (range 0-255).
None	To specify signatures independent of protocol, select None.
Timeout	Enter the number of seconds before an idle application flow is terminated (range 0-604800 seconds). A zero indicates that the default timeout of the application will be used. This value is used for protocols other than TCP and UDP in all cases and for TCP and UDP timeouts when the TCP timeout and UDP timeout are not specified.
TCP Timeout	Enter the number of seconds before an idle TCP application flow is terminated (range 0-604800 seconds). A zero indicates that the default timeout of the application will be used.
UDP Timeout	Enter the number of seconds before an idle UDP application flow is terminated (range 0-604800 seconds). A zero indicates that the default timeout of the application will be used.
Scanning	Select check boxes for the scanning types that you want to allow, based on security profiles (file types, data patterns, and viruses).

Table 100. New Application Settings (Continued)

Field	Description	
Signature Tab		
Signatures	Click Add to add a new signature, and specify the following information:	
	• Signature Name—Enter a name to identify the signature.	
	• Comment—Enter an optional description.	
	• Scope —Select whether to apply this signature only to the current transaction or to the full user session.	
	• Ordered Condition Match—Select if the order in which signature con- ditions are defined is important.	
	Specify conditions to define signatures:	
	• Add a condition by clicking Add AND Condition or Add OR Condi- tion . To add a condition within a group, select the group and then click Add Condition .	
	• Select an operator from Pattern Match and Equal To . When choosing a pattern match operator, specify the following:	
	 Context—Select from the available contexts. 	
	 Pattern—Specify a regular expression. See Table 104 for pattern rules for regular expressions. 	
	- Qualifier and Value—Optionally, add qualifier/value pairs.	
	 When choosing an equal to operator, specify the following, 	
	 Context—Select from unknown requests and responses for TCP or UDP. 	
	 Position—Select between the first four or second four bytes in the payload. 	
	 Mask—Specify a 4-byte hex value, for example, 0xfffff00. 	
	- Value—Specify a 4-byte hex value, for example, 0xaabbccdd.	
	• To move a condition within a group, select the condition and click the Move Up or Move Down arrow. To move a group, select the group and click the Move Up or Move Down arrow. You cannot move conditions from one group to another.	

Table 100. New Application Settings (Continued)



Note: It is not required to specify signatures for the application if the application is used only for application override rules.

To import an application, click **Import**. Browse to select the file, and select the target virtual system from the **Destination** drop-down list.

To export the application, select the check box for the application and click **Export**. Follow the prompts to save the file.

Custom Applications with Signatures

You can define custom applications with signatures. This section provides examples of how this can be done. Refer to the *PAN-OS Command Line Interface Reference Guide* for information on the show application command.

Example - Detect web traffic to a specified site

This example shows an application that detects web traffic going to *www.specifiedsite.com*.

Requests to the web site are of the following form:

```
GET /001/guest/
viewprofile.act?fa=25&tg=M&mg=F&searchType=zipcode&type=QUICK&pict=true&cont
ext=adrr&zip=94024&ta=34&sb=&item=0&pn=0 HTTP/1.1
Host: www.specifiedsite.com
```

```
User-Agent: Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US; rv:1.9.0.7)
Gecko/2009021910 Firefox/3.0.7 Accept: text/html,application/
xhtml+xml,application/xml;q=0.9,*/*;q=0.8 Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip,deflate Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
Keep-Alive: 300 Connection: keep-alive Referer: http://www.specifiedsite.com/
001/guest/
search.act?type=QUICK&pict=true&sb=&fa=25&ta=34&mg=F&tg=M&searchType=zipcode
&zip=94024&context=adrr&context=adrr Cookie:
JSESSIONID=A41B41A19B7533589D6E88190B7F0B3D.001; specifiedsite.com/
jumpcookie=445461346*google.com/search?q=lava+life&; locale=en_US;
campaign=1; imageNum=2; cfTag_LogSid=9327803497943a1237780204643;
___utma=69052556.1949878616336713500.1238193797.1238193797.1238193797.1;
__utmb=69052556.2.10.1238193797; __utmc=69052556;
___utmz=69052556.1238193797.1.1.utmcsr=(direct)|utmcmd=(none)
; utmv=69052556.gender%3Df; launch=1
```

The following signature can identify *specifiedsite* traffic if the host field is *www.specifiedsite.com*. username@hostname# **show application specifiedsite**

```
specifiedsite {
 category collaboration;
 subcategory social-networking;
 technology browser-based;
 decoder http;
 signature {
   s1 {
     and-condition {
       a1 {
         or-condition {
           01 {
             context http-reg-host-header;
             pattern www\.specifiedsite\.com;
}
}
```

Example - Detect a post to a specified blog

This example shows an application that detects blog posting activity on *www.specifiedblog.com*. In this example, it is not necessary to detect when somebody tries to read the blog, only to detect when an item is getting posted.

The post traffic request includes the following:

```
POST /wp-admin/post.php HTTP/1.1 Host: panqa100.specifiedblog.com
User-Agent: Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US; rv:1.9.0.7)
```

Gecko/2009021910 Firefox/3.0.7 Accept: text/html,application/ xhtml+xml,application/xml;q=0.9,*/*;q=0.8 Accept-Language: en-us,en;q=0.5 Accept-Encoding: gzip, deflate Accept-Charset: ISO-8859-1, utf-8; q=0.7,*; q=0.7 Keep-Alive: 300 Connection: keep-alive Referer: http:// panqa100.specifiedblog.com/wp-admin/post.php?action=edit&post=1 Cookie: utma=96763468.235424814.1238195613.1238195613.1238195613.1; ___utmb=96731468; utmc=96731468; utmz=96731468.1238195613.1.1.utmccn=(organic)|utmcsr=google|utmctr=blog+ho st|utmcmd=organic; wordpressuser bfbaae4493589d9f388265e737a177c8=panga100; wordpresspass bfbaae4493589d9f388265e737a177c8=c68a8c4eca4899017c58668eacc05 fc2 Content-Type: application/x-www-form-urlencoded Content-Length: 462 user_ID=1&action=editpost&post_author=1&post_ID=1&post_title=Hello+world%21& post category%5B%5D=1&advanced view=1&comment status=open&post password=&exc erpt=&content=Hello+world.%3Cbr+%2F%3E&use_instant_preview=1&post_pingback=1 &prev status=publish&submit=Save&referredby=http%3A%2F%2Fpanqa100.specifiedb log.com%2Fwp-admin%2F&post status=publish&trackback url=&post name=helloworld&post author override=1&mm=3&jj=27&aa=2009&hh=23&mn=14&ss=42&metakeyinp ut=&metavalue=HTTP/1.1

The host field includes the pattern *specifiedblog.com*. However, if a signature is written with that value in the host, it will match all traffic going to *specifiedblog.com*, including posting and viewing traffic. Therefore, it is necessary to look for more patterns.

One way to do this is to look for *post_title* and *post-author* patterns in the parameters of the post. The resulting signature detects postings to the web site:

```
username@hostname# show application specifiedblog blog posting
specifiedblog blog posting {
  category collaboration;
  subcategory web-posting;
  technology browser-based;
  decoder http;
  signature {
    s1 {
      and-condition {
        a1 {
          or-condition {
            01 {
              context http-req-host-header;
              pattern specifiedblog\.com;
              method POST;
            }
          }
        }
        a2 {
          or-condition {
            02 {
              context http-req-params;
              pattern post title;
              method POST;
            }
          }
        a3 {
          or-condition {
            03 {
              context http-reg-params;
              pattern post author;
              method POST;
}
}
```

Defining Application Groups

► Objects > Application Groups

To simplify the creation of security policies, applications requiring the same security settings can be combined into application groups. To define new applications, refer to "Defining Applications" on page 233.

	Table	101.	New	Application	Group
--	-------	------	-----	-------------	-------

Field	Description
Name	Enter a name that describes the application group (up to 31 characters). This name appears in the application list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Applications	Click Add and select applications, application filters, and/or other application groups to be included in this group.

Application Filters

▶ Objects > Application Filters

You can define application filters to simplify repeated searches. To define application filters to simplify repeated searches, click **Add** and enter a name for the filter.

In the upper area of the window, click an item that you want to use as a basis for filtering. For example, to restrict the list to the Networking category, click **networking**.

Search	٩	Custom Only	ar Filters	1244 matching applications
Category 🔺	Subcategory 🔺	Technology 🔺	Risk 🛥	Characteristic 🔺
241 business-systems	31 audio-streaming	476 browser-based	317 📘	499 Evasive
333 collaboration	11 auth-service	452 dient-server	242 2	394 Excessive Bandwidth
196 general-internet	15 database	197 network-protocol	294 3	264 Prone to Misuse
173 media	56 email	117 peer-to-peer	258 4	612 Transfers Files
299 networking	33 encrypted-tunnel		133 5	249 Tunnels Other Apps
2 unknown	19 erp-crm			266 Used by Malware
	154 file-sharing			733 Vulnerability
	46 gaming			779 Widely used

To filter on additional columns, select an entry in the columns to display check boxes. The filtering is successive: first category filters are applied, then sub category filters, then technology filters, then risk, filters, and finally characteristic filters.

For example, the next figure shows the result of choosing a category, sub category, and risk filter. In applying the first two filters, the **Technology** column is automatically restricted to the technologies that are consistent with the selected category and sub category, even though a technology filter has not been explicitly applied.

As you select options, the list of applications in the lower part of the page is automatically updated, as shown in the figure.

Search	٩	Custom Only	ar Filters	5 matching applications
Category 🔺	Subcategory 🔺	Technology 🔺	Risk 🔺	Characteristic 🔺
5 collaboration	56 email	5 browser-based	2 1	4 Transfers Files
	91 instant-messaging		1 2	1 Tunnels Other Apps
	30 internet-conferencing		2 3	4 Vulnerability
	5 social-business			2 Widely used
	64 social-networking			
	52 voip-video			
	35 web-posting			
Name	Category	Subcategory		Risk Technology
III blackboard	collaboration	social-business		1 browser-based
III dearspace	collaboration	social-business		3 browser-based
III projectplace	collaboration	social-business		2 browser-based
💷 rypple	collaboration	social-business		browser-based
III sharepoint (1 out of 6 shown)				
🦾 📑 sharepoint-base	collaboration	social-business		3 browser-based

Services

▶ Objects > Services

When you define security policies for specific applications, you can select one or more services to limit the port numbers the applications can use. The default service is **any**, which allows all TCP and UDP ports.

The HTTP and HTTPS services are predefined, but you can add additional service definitions. Services that are often assigned together can be combined into service groups to simplify the creation of security policies (refer to "Service Groups" on page 240).

Field	Description
Name	Enter the service name (up to 63 characters). This name appears in the services list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Description	Enter a description for the service (up to 255 characters).
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems.
Protocol	Select the protocol used by the service (TCP or UDP).
Destination Port	Enter the destination port number (0 to 65535) or range of port numbers (port1-port2) used by the service. Multiple ports or ranges must be separated by commas. The destination port is required.
Source Port	Enter the source port number (0 to 65535) or range of port numbers (port1-port2) used by the service. Multiple ports or ranges must be separated by commas. The source port is optional.

Service Groups

Objects > Services Groups

To simplify the creation of security policies, you can combine services that have the same security settings into service groups. To define new services, refer to "Services" on page 239.

Table 103. Service Group Settings

Field	Description
Name	Enter the service group name (up to 63 characters). This name appears in the services list when defining security policies. The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Service	Click Add to add services to the group. Select from the drop-down list, or click the Service button at the bottom of the drop-down list, and specify the settings. Refer to "Services" on page 239 for a description of the settings.

Data Patterns

Data pattern support allows you to specify categories of sensitive information that you may want to subject to filtering using data filtering security policies. For instructions on configuring data patterns, refer to "Dynamic Block Lists" on page 242.

When adding a new pattern (regular expression), the following general requirements apply:

- The pattern must have string of at least 7 bytes to match. It can contain more than 7 bytes, but not fewer.
- The string match is case-sensitive, as with most regular expression engines. Looking for "confidential" is different than looking for "Confidential" or "CONFIDENTIAL."

The regular expression syntax in PAN-OS is similar to traditional regular expression engines, but every engine is unique. The following table describes the syntax supported in PAN-OS.

Syntax	Description
•	Match any single character.
?	Match the preceding character or expression 0 or 1 time. The general expression MUST be inside a pair of parentheses.
	Example: (abc)?
*	Match the preceding character or expression 0 or more times. The general expression MUST be inside a pair of parentheses.
	Example: (abc)*
+	Match the preceding character or regular expression 1 or more times. The general expression MUST be inside a pair of parentheses.
	Example: (abc)+

Table 104. Pattern Rules

Syntax	Description
	Equivalent to "or".
	Example: ((bif) (scr) (exe)) matches "bif", "scr" or "exe". Note that the alternative substrings must be in parentheses.
-	Used to create range expressions.
	Example: [c-z] matches any character between c and z, inclusive.
[]	Match any.
	Example: [abz]: matches any of the characters a, b, or z.
^	Match any except.
	Example: [^abz] matches any character except a, b, or z.
{ }	Min/Max number of bytes.
	Example: {10-20} matches any string that is between 10 and 20 bytes. This must be directly in front of a fixed string, and only supports "-".
\	To perform a literal match on any one of the special characters above, it MUST be escaped by preceding them with a '\' (backslash).
&	& is a special character, so to look for the "&" in a string you must use "&" instead.

Table 104. Pattern Rules

Data Patterns Examples

The following are examples of valid custom patterns:

- .*((Confidential) | (CONFIDENTIAL))
 - Looks for the word "Confidential" or "CONFIDENTIAL" anywhere
 - ".*" at the beginning specifies to look anywhere in the stream
 - Does not match "confidential" (all lower case)
- .*((Proprietary & amp Confidential) | (Proprietary and Confidential))
 - Looks for either "Proprietary & Confidential" or "Proprietary and Confidential"
 - More precise than looking for "Confidential"
- .*(Press Release).*((Draft) | (DRAFT) | (draft))
 - Looks for "Press Release" followed by various forms of the word draft, which may indicate that the press release isn't ready to be sent outside the company
- .*(Trinidad)
 - Looks for a project code name, such as "Trinidad"

Custom URL Categories

Objects > Custom URL Categories

The custom URL categories feature allows you to create your own lists of URLs that can be selected in any URL filtering profile. Each custom category can be controlled independently and will have an action associated with it in each URL filtering profile (allow, block, continue, override, or alert).

URL entries can be added individually, or you can import a list of URLs. To do so, create a text file that contains the URLs to include, with one URL per line. Each URL can be in the format "www.example.com," and can contain * as a wildcard, such as "*.example.com." For additional information on wildcards, refer to the description of Block List in Table 90 on page 217.



Note: URL entries added to custom categories are case insensitive.

For instructions on setting up URL filtering profiles, refer to "URL Filtering Profiles" on page 217.

Field	Description
Name	Enter a name to identify the custom URL category (up to 31 characters). This name appears in the category list when defining URL filtering policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Description	Enter a description for the URL category (up to 255 characters).
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.
Sites	In the Sites area, click Add to enter a URL or click Import and browse to select the text file that contains the list of URLs.

Table 105. Custom URL Categories

Dynamic Block Lists

▶ Objects > Dynamic Block Lists

Use the **Dynamic Block Lists** page to create an address object based on an imported list of IP addresses. The source of the list must be a text file and must be located on a web server. You can set the **Repeat** option to automatically update the list on the device hourly, daily, weekly, or monthly. After creating a dynamic block list object, you can then use the address object in the source and destination fields for security policies. Each imported list can contain up to 5,000 IP addresses (IPv4 and/or IPv6), IP ranges, or subnets.

The list must contain one IP address, range, or subnet per line, for example:

"192.168.80.150/32" indicates one address, and "192.168.80.0/24" indicates all addresses from 192.168.80.0 through 192.168.80.255.

Example:

"2001:db8:123:1::1" or "2001:db8:123:1::/64"

IP Range:

To specify an address range, select **IP Range**, and enter a range of addresses. The format is: *ip_address_ip_address* where each address can be IPv4 or IPv6. Example: "2001:db8:123:1::1 - 2001:db8:123:1::22"

Field	Description
Name	Enter a name to identify the Dynamic Block List (up to 32 characters). This name will appear when selecting the source or destination in a policy.
Description	Enter a description for the block list (up to 255 characters).
Source	Enter an HTTP or HTTPS URL path that contains the text file. For example, http://1.1.1.1/myfile.txt.
Repeat	Specify the frequency in which the list should be imported. You can choose hourly, daily, weekly, or monthly. At the specified interval, the list will be imported into the configuration. A full commit is not needed for this type of update to occur.
Test Source URL	Test that the source URL or server path is available.

Table 106 Dynamic Block Lists

Custom Spyware and Vulnerability Signatures

This section describes the options available to create custom Spyware and Vulnerability signatures that can be used when creating custom vulnerability profiles.

- Objects > Custom Signatures > Data Patterns
- Objects > Custom Signatures > Spyware
- Objects > Custom Signatures > Vulnerability

Defining Data Patterns

▶ Objects > Custom Signatures > Data Patterns

Use the **Data Patterns** page to define the categories of sensitive information that you may want to subject to filtering using data filtering security policies. For information on defining data filtering profiles, refer to "Data Filtering Profiles" on page 223.

Field	Description
Name	Enter the data pattern name (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Description	Enter a description for the data pattern (up to 255 characters).

Table 107. Data Pattern Settings

Field	Description
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.
Weight	Enter weights for pre-specified pattern types. The weight is a number between 1 and 255. Alert and Block thresholds specified in the Data Filtering Profile are a function of this weight.
	• CC#—Specify a weight for the credit card field (range 0-255).
	• SSN# —Specify a weight for the social security number field, where the field includes dashes, such as 123-45-6789 (range 0-255, 255 is highest weight).
	• SSN# (without dash) —Specify a weight for the social security number field, where the entry is made without dashes, such as 123456789 (range 0-255, 255 is highest weight).
Custom Patterns	The pre-defined patterns include credit card number and social security number (with and without dashes).
	Click Add to add a new pattern. Specify a name for the pattern, enter the regular expression that defines the pattern, and enter a weight to assign to the pattern. Add additional patterns as needed.

Table 107. Data Pattern Settings (Continued)

Defining Spyware and Vulnerability Signatures

Objects > Custom Signatures > Spyware & Vulnerabilities

The firewall supports the ability to create custom spyware and vulnerability signatures using the firewall threat engine. You can write custom regular expression patterns to identify spyware phone home communication or vulnerability exploits. The resulting spyware and vulnerability patterns become available for use in any custom vulnerability profiles. The firewall looks for the custom-defined patterns in network traffic and takes the specified action for the vulnerability exploit. Support is provided for creation of custom signatures using HTTP, SMTP, IMAP, FTP, POP3, SMB, MSSQL, MSRPC, RTSP, SSH, SSL, Telnet, Unknown-TCP, and Unknown-UDP.

You can optionally include a time attribute when defining custom signatures by specifying a threshold per interval for triggering possible actions in response to an attack. Action is taken only after the threshold is reached.

Use the **Custom Signatures** page to define signatures for vulnerability profiles.

Field	Description
Configuration Tab	
Threat ID	Enter a numeric identifier for the configuration. For spyware signatures, the range is 15000-18000; for vulnerability signatures the range is 41000-45000.
Name	Specify the threat name.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.
Comment	Enter an optional comment.

Table 108. Custom Signatures - Vulnerability and Spyware

Field	Description
Severity	Assign a level that indicates the seriousness of the threat.
Default Action	Assign the default action to take if the threat conditions are met:
	• Alert—Generate an alert.
	• Drop Packets—Do not allow packets through.
	• Reset Both —Reset the client and server.
	• Reset Client—Reset the client.
	• Reset Server —Reset the server.
	• Block IP —Block traffic for a specified period of time. Choose whether to block traffic for the source only or source and destination, and enter the duration (seconds).
Direction	Indicate whether the threat is assessed from the client to server, server to client, or both.
Affected System	Indicate whether the threat involves the client, server, either, or both. Applies to vulnerability signatures, but not spyware signatures.
CVE	Specify the common vulnerability enumeration (CVE) as an external reference for additional background and analysis.
Vendor	Specify the vendor identifier for the vulnerability as an external reference for additional background and analysis.
Bugtraq	Specify the bugtraq (similar to CVE) as an external reference for additional background and analysis.
Reference	Add any links to additional analysis or background information. The information is shown when a user clicks on the threat from the ACC, logs, or vulnerability profile.
Signatures Tab	
Standard Signature	Select the Standard radio button and then click Add to add a new signature. Specify the following information:
	• Standard —Enter a name to identify the signature.
	• Comment —Enter an optional description.
	• Ordered Condition Match—Select if the order in which signature con- ditions are defined is important.
	• Scope —Select whether to apply this signature only to the current trans- action or to the full user session.
	Specify conditions to define signatures:
	• Add a condition by clicking Add AND Condition or Add OR Condi- tion . To add a condition within a group, select the group and then click Add Condition . Select from the Method and Context drop-down lists. Specify a regular expression in the Pattern field. Add additional pat- terns as needed.
	• To move a condition within a group, select the condition and click the Move Up or Move Down arrow. To move a group, select the group and click the Move Up or Move Down arrow. You cannot move conditions from one group to another.

Table 108. Custom Signatures - Vulnerability and Spyware (Continued)

Field	Description
Combination Signature	Select the Combination radio button. In the area above the subtabs, specify the following information:
	On the Combination Signatures subtab, specify conditions to define signatures:
	• Add a condition by clicking Add AND Condition or Add OR Condi- tion . To add a condition within a group, select the group and then click Add Condition . Select from the Method and Context drop-down lists. Specify a regular expression in the Pattern field. Add additional pat- terns as needed.
	• To move a condition within a group, select the condition and click the Move Up or Move Down arrow. To move a group, select the group and click the Move Up or Move Down arrow. You cannot move conditions from one group to another.
	On the Time Attribute subtab, specify the following information:
	• Number of Hits—Specify the threshold that will trigger any policy- based action as a number of hits (1-1000) in a specified number of seconds (1-3600).
	• Aggregation Criteria—Specify whether the hits are tracked by source IP address, destination IP address, or a combination of source and destination IP addresses.

Table 108. Custom Signatures - Vulnerability and Spyware (Continued)

Security Profile Groups

Objects > Security Profile Groups

The firewall supports the ability to create security profile groups, which specify sets of security profiles that can be treated as a unit and then added to security policies. For example, you can create a "threats" security profile group that includes profiles for antivirus, anti-spyware, and vulnerability and then create a security policy that includes the "threats" profile.

Antivirus, anti-spyware, vulnerability protection, URL filtering, and file blocking profiles that are often assigned together can be combined into profile groups to simplify the creation of security policies.

To define new security profiles, refer to "Defining Security Policies" on page 187.

Field	Description
Name	Enter the profile group name (up to 31 characters). This name appears in the profiles list when defining security policies. The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow the profile to be shared by all virtual systems.

Table 109. Security Profile Group Settings

Field	Description
Profiles	Select an antivirus, anti-spyware, vulnerability protection, URL filtering, and/or file blocking profile to be included in this group. Data filtering profiles can also be specified in security profile groups. Refer to "Data Filtering Profiles" on page 223.

Table 109. Security Profile Group Settings (Continued)

Log Forwarding

Objects > Log Forwarding

Each security policy can specify a log forwarding profile that determines whether traffic and threat log entries are logged remotely with Panorama, and/or sent as SNMP traps, syslog messages, or email notifications. By default, only local logging is performed.

Traffic logs record information about each traffic flow, and threat logs record the threats or problems with the network traffic, such as virus or spyware detection. Note that the antivirus, anti-spyware, and vulnerability protection profiles associated with each rule determine which threats are logged (locally or remotely). To apply logging profiles to security policies, refer to "Security Policies" on page 187.

Field	Description
Name	Enter a profile name (up to 31 characters). This name appears in the list of log forwarding profiles when defining security policies. The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems.
Traffic Settings	
Panorama	Select the check box to enable sending traffic log entries to the Panorama centralized management system. To define the Panorama server address, refer to "Defining Management Settings" on page 30.
SNMP Trap Email Syslog	Select the SNMP, syslog, and/or email settings that specify additional destinations where the traffic log entries are sent. To define new destinations, refer to:
	"Configuring SNMP Trap Destinations" on page 75.
	• "You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in Device > Server Profiles > Syslog. Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type." on page 77
	 "Configuring Syslog Servers" on page 76

Table 110. Log Forwarding Profile Settings

Field	Description
Threat Log Settings	
Panorama	Click the check box for each severity level of the threat log entries to be sent to Panorama. The severity levels are:
	• Critical—Very serious attacks detected by the threat security engine.
	• High—Major attacks detected by the threat security engine.
	 Medium—Minor attacks detected by the threat security engine, including URL blocking.
	• Low—Warning-level attacks detected by the threat security engine.
	• Informational —All other events not covered by the other severity levels, including informational attack object matches.
SNMP Trap Email Syslog	Under each severity level, select the SNMP, syslog, and/or email settings that specify additional destinations where the threat log entries are sent.

Table 110. Log Forwarding Profile Settings (Continued)

Decryption Profiles

► Objects > Decryption Profiles

Decryption profiles enable you to block and control specific aspects of the SSL forward proxy, SSL inbound inspection, and SSH traffic. After you create a decryption profile, you can then apply that profile to a decryption policy.

You can also control the trusted CAs that your device trusts, for more information, refer to "Default Trusted Certificate Authorities" on page 88.

Table 111. Decryption Profile Settings

Field	Description
SSL Forward Proxy Tab	
Server Certificate Checks	Select options to control server certificates.
Block sessions with expired certificates	Terminate the SSL connection if the server certificate is expired. This will prevent a user from being able to accept an expired certificate and continuing with an SSL session.
Block sessions with untrusted issuers	Terminate the SSL session if the server certificate issuer is untrusted.
Restrict certificate extensions	Limits the certificate extensions used in the dynamic server certificate to key usage and extended key usage.
	Details—Displays details on the values used for key usage and extended key usage.
Unsupported Mode Checks	Select options to control unsupported SSL applications.
Block sessions with unsupported version	Terminate sessions if the "client hello" message is not supported by PAN-OS. The SSL versions supported by PAN-OS are: SSLv3, TLS1.0, and TLS1.1.

Field	Description
Block sessions with unsupported cipher suites	Terminate the session if the cipher suite specified in the SSL handshake if it is not supported by PAN-OS.
Block sessions with client authentication	Terminate sessions with client authentication for SSL forward proxy traffic.
Failure Checks	Select the action to take if system resources are not available to process decryption.
Block sessions if resources not available	Terminate sessions if system resources are not available to process decryption.
	Note: For unsupported modes and failure modes, the session information is cached for 12 hours, so future sessions between the same hosts and server pair are not decrypted. Use the check boxes to block those sessions instead.
SSL Inbound Inspection Tab	
Unsupported Mode Checks	Selection options to control sessions if unsupported modes are detected in SSL traffic.
Block sessions with unsupported versions	Terminate sessions if the "client hello" message is not supported by PAN-OS. The SSL versions supported by PAN-OS are: SSLv3, TLS1.0, and TLS1.1.
Block sessions with unsupported cipher suites	Terminate the session if the cipher suite used is not supported by PAN-OS.
Failure Checks	Select the action to take if system resources are not available.
Block sessions if resources not available	Terminate sessions if system resources are not available to process decryption.
SSH Tab	
Unsupported Mode Checks	Selection options to control sessions if unsupported modes are detected in SSH traffic. Supported SSH version is SSH version 2.
Block sessions with unsupported versions	Terminate sessions if the "client hello" message is not supported by PAN-OS.
Block sessions with unsupported algorithms	Terminate sessions if the algorithm specified by the client or server is not supported by PAN-OS.
Failure Checks	Select actions to take if SSH application errors occur and if system resources are not available.

Table 111. Decryption Profile Settings

Field	Description
Block sessions on SSH errors	Terminate sessions if SSH errors occur.
Block sessions if resources not available	Terminate sessions if system resources are not available to process decryption.

Table 111. Decryption Profile Settings

Schedules

► Objects > Schedules

By default, each security policy applies to all dates and times. To limit a security policy to specific times, you can define schedules, and then apply them to the appropriate policies. For each schedule, you can specify a fixed date and time range or a recurring daily or weekly schedule. To apply schedules to security policies, refer to "Security Policies" on page 187.



Note: When a security policy is invoked by a defined schedule, only new sessions are affected by the applied security policy. Existing sessions are not affected by the scheduled policy.

Field	Description
Name	Enter a schedule name (up to 31 characters). This name appears in the schedule list when defining security policies. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	If the device is in Multiple Virtual System Mode, select this check box to allow sharing by all virtual systems.
Recurrence	Select the type of schedule (Daily, Weekly, or Non-Recurring).
Daily	Click Add and specify a start and end time in 24-hour format (HH:MM).
Weekly	Click Add , select a day of the week, and specify the start and end time in 24-hour format (HH:MM).
Non-recurring	Click Add and specify a start and end date and time.

Table 112. Schedule Settings

Other Policy Objects

Chapter 6 Reports and Logs

This chapter describes how to view the reports and logs provided with the firewall:

- "Using the Dashboard" in the next section
- "Using the Application Command Center" on page 253
- "Using App-Scope" on page 256
- "Viewing the Logs" on page 264
- "Working with Botnet Reports" on page 267
- "Managing PDF Summary Reports" on page 270
- "Managing User Activity Reports" on page 272
- "Managing Report Groups" on page 272
- "Scheduling Reports for Email Delivery" on page 273
- "Viewing Reports" on page 273
- "Generating Custom Reports" on page 274
- "Identifying Unknown Applications and Taking Action" on page 276
- "Taking Packet Captures" on page 278



Note: Most of the reports in this section support optional selection of a virtual system from the drop-down list at the top of page.

Using the Dashboard

▶ Dashboard

The **Dashboard** page displays general device information, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available charts are displayed by default, but each user can remove and add individual charts, as needed.

Click **Refresh** to update the Dashboard. To change the automatic refresh interval, select an interval from the drop-down list (1 min, 2 mins, 5 mins, or Manual). To add a Widget to the Dashboard, click the Widget drop-down, select a category and then the widget name. To delete a widget, click 🔀 in the title bar.

Review the following information in each chart.

Chart	Description
Top Applications	Displays the applications with the most sessions. The block size indicates the relative number of sessions (mouse-over the block to view the number), and the color indicates the security risk—from green (lowest) to red (highest). Click an application to view its application profile.
Top High Risk Applications	Similar to Top Applications, except that it displays the highest-risk applications with the most sessions.
General Information	Displays the device name, model, PAN-OS software version, the application, threat, and URL filtering definition versions, the current date and time, and the length of time since the last restart.
Interface Status	Indicates whether each interface is up (green), down (red), or in an unknown state (gray).
Threat Logs	Displays the threat ID, application, and date and time for the last 10 entries in the Threat log. The threat ID is a malware description or URL that violates the URL filtering profile.
Config Logs	Displays the administrator user name, client (Web or CLI), and date and time for the last 10 entries in the Configuration log.
Data Filtering Logs	Displays the description and date and time for the last 60 minutes in the Data Filtering log.
URL Filtering Logs	Displays the description and date and time for the last 60 minutes in the URL Filtering log.
System Logs	Displays the description and date and time for the last 10 entries in the System log. Note that a "Config installed" entry indicates configuration changes were committed successfully.
System Resources	Displays the Management CPU usage, Data Plane usage, and the Session Count, which displays the number of sessions established through the firewall.
Logged In Admins	Displays the source IP address, session type (Web or CLI), and session start time for each administrator who is currently logged in.
ACC Risk Factor	Displays the average risk factor (1 to 5) for the network traffic processed over the past week. Higher values indicate higher risk.

Table 113. Dashboard Charts
Chart	Description
High Availability	If high availability (HA) is enabled, indicates the HA status of the local and peer device—green (active), yellow (passive), or black (other). For more information about HA, refer to "Enabling HA on the Firewall" on page 101.
Locks	Shows configuration locks taken by administrators.

Table 113. Dashboard Charts (Continued)

Using the Application Command Center

► ACC

The **Application Command Center (ACC)** page displays the overall risk level for your network traffic, the risk levels and number of threats detected for the most active and highest-risk applications on your network, and the number of threats detected from the busiest application categories and from all applications at each risk level. The ACC can be viewed for the past hour, day, week, month, or any custom-defined time frame.

Risk levels (1=lowest to 5=highest) indicate the application's relative security risk based on criteria such as whether the application can share files, is prone to misuse, or tries to evade firewalls.

To view the Application Command Center:

- 1. Under the **ACC** tab, change one or more of the following settings at the top of the page, and click **Go**:
 - a. Select a virtual system, if virtual systems are defined.
 - b. Select a time period from the **Time** drop-down list. The default is Last Hour.
 - c. Select a sorting method from the **Sort By** drop-down list. You can sort the charts in descending order by number of sessions, bytes, or threats. The default is by number of sessions.
 - d. For the selected sorting method, select the top number of applications and application categories shown in each chart from the **Top** drop-down list.

		Virtual S ₁	vstem All	Time Last Hour	Sort By Session	s 🔻 Top 25	▼ → +	
lication Command (enter					R, R) () () ()	2 3 4
	Applio	ation					Applications	~
		Risk	Application Name	Sessions		Bytes		Threats
	1	2	dns	2.1K	421.9 K			0 🛙
	2	1	insufficient-data	1.5K	108.2 K			0
	3	2	google-safebrowsing	4	28.1 K			0
	4	2	ipsec-esp	3	450	1		0
	5	2	ntp	2	180	1		0
	6	5	rss	2	107.1 K			0 🛙
	URL F	iltering					URL Categor	ies 💌
		Catego	ory		Sessions			Bytes
	1	search	-engines		4		28.	1K
	2	news			2		106.	4K

Figure 22. Application Command Center Page

2. To open log pages associated with the information on the page, use the log links in the upper-right corner of the page, as shown here. The context for the logs matches the information on the page.



3. To filter the list, click an item in one of the columns, this will add that item to the filter bar located above the log column names. After adding the desired filters, click the Apply Filter icon.

🔍 (rece	(receive_time in last-hour) and (zone.src eq I3-vlan-trust) and (addr.src in 192.168.2.10)						🛱 🖳
	Receive Time	Туре	From Zone	To Zone	Source	Source User	Destinatio
Þ	09/05 16:44:14	end	13-vlan-trust	13-untrust	192.168.2.10		10.0. 🔺
P	09/05 16:44:14	end	I3-vlan-trust	13-untrust	192.168.2.10		10.0.) =

- 4. Choose a view from the drop-down list for the area of interest, as described in the following table.
- 5. Use the drop-down lists for Applications, URL Categories, Threats, Content/File Types, and HIP Objects.

Chart	Description
Application	Displays information organized according to the menu selection. Information includes the number of sessions, bytes transmitted and received, number of threats, application category, application subcategories, application technology, and risk level, as applicable.
	• Applications
	• High risk applications
	Categories
	Sub Categories
	• Technology
	• Risk
URL Filtering	Displays information organized according to the menu selection. Information includes the URL, URL category, repeat count (number of times access was attempted, as applicable).
	URL Categories
	• URLs
	Blocked URL Categories
	Blocked URLs

Table 114. Application Command Center Charts

Chart	Description
Threat Prevention	Displays information organized according to the menu selection. Information includes threat ID, count (number of occurrences), number of sessions, and subtype (such as vulnerability), as applicable.
	• Threats
	• Types
	• Spyware
	• Spyware Phone Home
	• Spyware Downloads
	• Vulnerability
	• Virus
Data Filtering	• Content/File Types
	• Types
	• File Names
HIP Matches	• HIP Objects
	• HIP Profiles

Table 114. Application Command Center Charts (Continued)

6. To view additional details, click any of the links. A details page opens to show information about the item at the top and additional lists for related items.

	pplication Information							
	Name: Description:	web-browsing Web Browsing is us Wide Web. Its origi	ing Hypertext Transfer I inal purpose was to prov	Protocol (HTTP), which is vide a way to publish and	s a method used d retrieve HTML	to transfer or co pages.	onvey infor	mation on the World
	Standard Ports: Capable of File Transfer: Used by Malware: Excessive Bandwidth Use: Evasive: Tunnels Other Applications: Additional Information:	: tcp/80 : yes : no : no : yes : Wikipedia Google	Yahoo!	Has Known Vu Pror Session Timeou TCP Timeou UDP Timeou	Category: Subcategory: Technology: Risk: Widely Used: ulnerabilities: ne to Misuse: ut (seconds): ut (seconds): ut (seconds):	general-interni internet-utility browser-based ves yes yes no	et J	
Т	pp Applications			Sest	sions			Bytes
To 1 To	Pp Applications Risk Application T web-browsing Pp Sources			Sess	sions		11.1M	Bytes

Figure 23. Application Command Center Drill Down Page

Using App-Scope

► Monitor > App Scope

The App-Scope reports introduce a visibility and analysis tools to help pinpoint problematic behavior, helping you understand the following aspects of your network:

- Changes in application usage and user activity
- Users and applications that take up most of the network bandwidth
- Network threats

With the App-Scope reports, you can quickly see if any behavior is unusual or unexpected. Each report provides a dynamic, user-customizable window into the network. The reports include options to select the data and ranges to display.

To view the reports, click the report name under **App-Scope** on the left side of the page in the **Monitor** tab. Select one of the report types lists below. Report options are available from the drop-down lists at the top and bottom of some of the pages.

Chart	Description
Summary	"Summary Report" on page 257
Change Monitor	"Change Monitor Report" on page 258
Threat Monitor	"Threat Monitor Report" on page 259
Threat Map	"Threat Monitor Report" on page 259
Network Monitor	"Network Monitor Report" on page 261
Traffic Map	"Traffic Map Report" on page 263

Table 115. Application Command Center Charts

Summary Report

The Summary report (Figure 24) displays charts for the top five gainers, losers, and bandwidth consuming applications, application categories, users, and sources.



Figure 24. App-Scope Summary Report

Change Monitor Report

The Change Monitor report (Figure 25) displays changes over a specified time period. For example, Figure 25 displays the top applications that gained in use over the last hour as compared with the last 24-hour period. The top applications are determined by session count and sorted by per cent.



Figure 25. App-Scope Change Monitor Report

This report contains the following buttons and options.

Table 116. Change Monitor Report Options

Item	Description
Top Bar	
<u>∎</u> Top 10 ▼	Determines the number of records with the highest measurement included in the chart.
P Application -	Determines the type of item reported: Application, Application Category, Source, or Destination.
Z Gainers	Displays measurements of items that have increased over the measured period.
Sa Losers	Displays measurements of items that have decreased over the measured period.
O New	Displays measurements of items that were added over the measure period.

ltem	Description
ع <u>ع</u> Z Dropped	Displays measurements of items that were discontinued over the measure period.
Filter 🖄 None 🔻	Applies a filter to display only the selected item. None displays all entries.
010	Determines whether to display session or byte information.
Sort: 😒 #	Determines whether to sort entries by percentage or raw growth.
Bottom Bar	
Compare last hour v to the same period ending 24 hours v ago	Specifies the period over which the change measurements are taken.

Table 116. Change Monitor Report Options (Continued)

Threat Monitor Report

The Threat Monitor report (Figure 26) displays a count of the top threats over the selected time period. For example, Figure 26 shows the top 10 threat types for the past 6 hours.



Figure 26. App-Scope Threat Monitor Report

Each threat type is color-coded as indicated in the legend below the chart. This report contains the following buttons and options.

Button	Description
Top Bar	
☐ Top 10 ▼	Determines the number of records with the highest measurement included in the chart.
P Threats •	Determines the type of item measured: Threat, Threat Category, Source, or Destination.
Filter 😺 🔯 🦻 😻 🖷	Applies a filter to display only the selected type of items.
	Determines whether the information is presented in a stacked column chart or a stacked area chart.
Bottom Bar	
Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days	Specifies the period over which the measurements are taken.

Table 117. Threat Monitor Report Buttons

Threat Map Report

The Threat Map report (Figure 27) shows a geographical view of threats, including severity.



Figure 27. App-Scope Threat Monitor Report

Each threat type is color-coded as indicated in the legend below the chart. Click a country on the map to zoom in. Click the **Zoom Out** button in the lower right corner of the screen to zoom out. This report contains the following buttons and options.

Button	Description
Top Bar	
🛄 Тор 10 🔻	Determines the number of records with the highest measurement included in the chart.
Incoming threats	Displays incoming threats.
Outgoing threats	Displays outgoing threats.
Filter 😼 🖏 🧔 😈 🖷	Applies a filter to display only the selected type of items.
Bottom Bar	
Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days	Indicates the period over which the measurements are taken.

Table 118. Threat Map Report Buttons

Network Monitor Report

The Network Monitor report (Figure 28) displays the bandwidth dedicated to different network functions over the specified period of time. Each network function is color-coded as indicated in the legend below the chart. For example, Figure 28 shows application bandwidth for the past 7 days based on session information.



Figure 28. App-Scope Network Monitor Report

The report contains the following buttons and options.

Button	Description
Top Bar	
Top 10 ▼	Determines the number of records with the highest measurement included in the chart.
Application 🔻	Determines the type of item reported: Application, Application Category, Source, or Destination.
Filter 🖄 None 🔻	Applies a filter to display only the selected item. None displays all entries.
010	Determines whether to display session or byte information.
	Determines whether the information is presented in a stacked column chart or a stacked area chart.
Bottom Bar	
Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days	Indicates the period over which the change measurements are taken.

Table 119. Network Monitor Report Buttons

Traffic Map Report

The Traffic Map report (Figure 29) shows a geographical view of traffic flows according to sessions or flows.



Figure 29. App-Scope Traffic Monitor Report

Each traffic type is color-coded as indicated in the legend below the chart. This report contains the following buttons and options.

Table 120. Threat Map Report Buttons

Button	Description
Top Bar	
Top 10 ▼	Determines the number of records with the highest measurement included in the chart.
Incoming threats	Displays incoming threats.
Outgoing threats	Displays outgoing threats.
010 101	Determines whether to display session or byte information.

	()
Button	Description
Bottom Bar	
Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days	Indicates the period over which the change measurements are taken.

Table 120. Threat Map Report Buttons (Continued)

Viewing the Logs

▶ Monitor > Logs

The firewall maintains logs for WildFire, configurations, system, alarms, traffic flows, threats, URL filtering, data filtering, and Host Information Profile (HIP) matches. You can view the current logs at any time. To locate specific entries, you can apply filters to most of the log fields.



Note: The firewall displays the information in logs so that role-based administration permissions are respected. When you display logs, only the information that you have permission to see is included. For information on administrator permissions, refer to "Defining Administrator Roles" on page 58.

To view the logs, click the log types on the left side of the page in the **Monitor** tab. Each log page has a filter area at the top of the page.

```
( zone.dst eq tapzone ) and ( app eq unknown-udp )
```

i 🔿 🗙 🕂 📴 🚔 🖄

Use the filter area as follows:

• Click any of the underlined links in the log listing to add that item as a log filter option. For example, if you click the **Host** link in the log entry for 10.0.0.252 and **Web Browsing** in both items are added, and the search will find entries that match both (AND search).



Note: To filter information for a user group, add the query (user in 'domain/user-group'). For example, to generate a list of the activities for the marketing group in Company ABC, you would query (user in 'company_abc/marketing').

• To define other search criteria, click the **Add Log Filter** icon. Select the type of search (and/or), the attribute to include in the search, the matching operator, and the values for the match, if appropriate. Click **Add** to add the criterion to the filter area on the log page, and then click **Close** to close the pop-up window. Click the **Apply Filter** icon to display the filtered list.



Note: You can combine filter expressions added on the log page with those that you define in the Expression pop-up window. Each is added as an entry on the Filter line on the log page.

If you set the "in" Received Time filter to **Last 60 seconds**, some of the page links on the log viewer may not show results because the number of pages may grow or shrink due to the dynamic nature of the selected time.

- To clear filters and redisplay the unfiltered list, click the **Clear Filter** button.
- To save your selections as a new filter, click the **Save Filter** button, enter a name for the filter, and click **OK**.
- To export the current log listing (as shown on the page, including any applied filters) click the **Save Filter** button. Select whether to open the file or save it to disk, and select the check box if you want to always use the same option. Click **OK**.

To change the automatic refresh interval, select an interval from the drop-down list (1 min, 30 seconds, 10 seconds, or Manual). To change the number of log entries per page, select the number of rows from the **Rows** drop-down list.

Log entries are retrieved in blocks of 10 pages. Use the paging controls at the bottom of the page to navigate through the log list. Select the **Resolve Hostname** check box to begin resolving external IP addresses to domain names.

Log Details								(
General						Time			
Session ID	140177		IP Pr	otocol u	ıdp	Generate	e Time 20)11/01/03 1	0:44:13
Туре	end		Log /	Action l	og-all	Star	t Time 20	011/01/03 1	0:43:21
Action	allow			Bytes 1	5,090	Receive	e Time 20	011/01/03 1	0:44:13
Application	unknown-	udp	Repeat	Count 1	L	Elapsed Time	(sec) 22	2	
Rule	Monitor A		Pa	ackets 🛛 🕄	36	Misc			
Category	any						Captive	Portal	R
Virtual System	Ho Vsys						D	L L	
Device	0003C101	1573					Proxy Tran	saction	2
Config Version	1						Dec	rypted	Ch
Source		Destinatio	n						<u></u>
Source User		Destina	ation User				Packet C	apture	R.
Source address 1	92.168.1.2	Destinatio	n address	192.16	58.1.1		Client to	Server	
Source Port 1	812	Destin	ation Port	2048			Cherre co		22
Source Zone ta	apzone	Destina	tion Zone	tapzor	ne		Server to	o Client	R.
Inbound Interface e	thernet 1/3	Outbound	Interface	ethern	et1/3			L	
									_
Related Logs									
Receive Time Log	Type Ap	plication	Action	Rule	Byte	s Packets	Severity	Category	URL
01/03 10:44:13 traffic	end unk	nown-udp	allow	Monitor A	15,09	0 36			
	_	_					_	_	
								Cla	
								CIO	se

To display additional details, click the spyglass icon 🦻 for an entry.

Figure 30. Log Entry Details

If the source or destination has an IP address to name mapping defined in the **Addresses** page, the name is presented instead of the IP address. To view the associated IP address, move your cursor over the name.

Review the following information in each log.

Table 1	21.	Log	Descriptions
---------	-----	-----	--------------

Chart	Description
Traffic	Displays an entry for the start and end of each session. Each entry includes the date and time, the source and destination zones, addresses, and ports, the application name, the security rule name applied to the flow, the rule action (allow, deny, or drop), the ingress and egress interface, and the number of bytes.
	Click next to an entry to view additional details about the session, such as whether an ICMP entry aggregates multiple sessions between the same source and destination (the Count value will be greater than one).
	Note that the Type column indicates whether the entry is for the start or end of the session, or whether the session was denied or dropped. A "drop" indicates that the security rule that blocked the traffic specified "any" application, while a "deny" indicates the rule identified a specific application.
	If traffic is dropped before the application is identified, such as when a rule drops all traffic for a specific service, the application is shown as "not-applicable".
Threat	Displays an entry for each security alarm generated by the firewall. Each entry includes the date and time, a threat name or URL, the source and destination zones, addresses, and ports, the application name, and the alarm action (allow or block) and severity.
	Click next to an entry to view additional details about the threat, such as whether the entry aggregates multiple threats of the same type between the same source and destination (the Count value will be greater than one).
	Note that the Type column indicates the type of threat, such as "virus" or "spyware." The Name column is the threat description or URL, and the Category column is the threat category (such as "keylogger") or URL category.
	If local packet captures are enabled, click 📮 next to an entry to access the captured packets, as in the following figure. To enable local packet captures, refer to the subsections under "Security Profiles" on page 211.
URL Filtering	Displays logs for URL filters, which block access to specific web sites and web site categories or generate an alert when a proscribed web site is accessed. Refer to "URL Filtering Profiles" on page 217 for information on defining URL filtering profiles.
WildFire	Displays logs for files that are uploaded and analyzed by the WildFire server, log data is sent back to the device after analysis, along with the analysis results.
	A subscription is required for this feature. If you do not have a subscription, you can use the WildFire Portal to view log information. Refer to "Using the WildFire Portal" on page 425.
Data Filtering	Displays logs for the security policies that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall. Refer to "Data Filtering Profiles" on page 223 for information on defining data filtering profiles.
	To configure password protection for access the details for a log entry, click the icon. Enter the password and click OK . Refer to "Defining Custom Response Pages" on page 115 for instructions on changing or deleting the data protection password.
	<i>Note:</i> The system prompts you to enter the password only once per session.

Chart	Description
Configuration	Displays an entry for each configuration change. Each entry includes the date and time, the administrator user name, the IP address from where the change was made, the type of client (Web or CLI), the type of command executed, whether the command succeeded or failed, the configuration path, and the values before and after the change.
System	Displays an entry for each system event. Each entry includes the date and time, the event severity, and an event description.
HIP Match	Displays information about security policies that apply to GlobalProtect clients. For more information, refer to "Overview" on page 335.
Alarms	The alarms log records detailed information on alarms that are generated by the system. The information in this log is also reported in the Alarms window. Refer to "Viewing Alarms" on page 85.

Table 121. Log Descriptions (Continued)

Viewing Session Information

Monitor > Session Browser

Open the **Session Browser** page to browse and filter current running sessions on the firewall. For information on filtering options for this page, refer to "Viewing the Logs" on page 264.

Working with Botnet Reports

The botnet report feature allows you to use behavior-based mechanisms to identify potential botnet-infected hosts in the network. Using network, threat, URL, and data filtering logs, the firewall evaluates threats based on criteria that include visits to malware sites and dynamic DNS sites, visits to recently registered domains (within the last 30 days), unknown application usage, and the existence of Internet Relay Chat (IRC) traffic.

After correlating and identifying hosts that match infected botnet behavior, the firewall assigns each potentially infected host a confidence score of 1 to 5 to indicate the likelihood of botnet infection (1 indicates the lowest and 5 the highest likelihood of infection). Because behavior-based detection mechanisms require correlating traffic across multiple logs over a period of 24 hours, the firewall generates a report every 24 hours that contains a sorted list of hosts based on confidence level.

Configuring the Botnet Report

Monitor > Botnet

Use these settings to specify types of suspicious traffic (traffic that may indicate botnet activity). To configure the settings, click the **Configuration** button on the right side of the **Botnet** page.

Field	Description
HTTP Traffic	Select the Enable check box for the events that you want to include in the reports:
	 Malware URL visit—Identifies users communicating with known malware URLs based on malware and botnet URL filtering categories.
	• Use of dynamic DNS—Looks for dynamic DNS query traffic that could indicate botnet communication.
	• Browsing to IP domains —Identifies users that browse to IP domains instead of URLs.
	• Browsing to recently registered domains —Looks for traffic to domains that have been registered within the past 30 days.
	• Executable files from unknown sites—Identifies executable files downloaded from unknown URLs.
Unknown Applications	Select the check boxes to mark unknown TCP or unknown UDP applications as suspicious, and specify the following information:
	• Sessions Per Hour—Number of application sessions per hour that are associated with the unknown application.
	• Destinations Per Hour —Number of destinations per hour that are associated with the unknown application.
	Minimum Bytes—Minimum payload size
	• Maximum Bytes—Maximum payload size.
IRC	Select the check box to include IRC servers as suspicious.

Table 122. Botnet Configuration Settings

Managing Botnet Reports

Monitor > Botnet > Report Setting

You can specify report queries and then generate and view botnet analysis reports. The reports are based on botnet configuration settings (refer to "Configuring the Botnet Report" on page 268). You can include or exclude source or destination IP addresses, users, zones, interfaces, regions, or countries.

Scheduled reports run once per day. You can also generate and display reports on demand by clicking **Run Now** in the window where you define the report queries. The generated report is displayed on the **Botnet** page.

To manage botnet reports, click the **Report Setting** button on the right side of the **Botnet** page. To export a report, select the report and click **Export to PDF** or **Export to CSV**.

Field	Description
Test Run Time Frame	Select the time interval for the report (last 24 hours or last calendar day).
# Rows	Specify the number of rows in the report.
Scheduled	Select the check box to run the report on a daily basis. To run the report manually, click Run Now at the top of the Botnet Report window.
Query	Construct the report query by specifying the following, and then clicking Add to add the configured expression to the query. Repeat as needed to construct the complete query:
	• Connector—Specify a logical connector (AND/OR).
	• Attribute—Specify the source or destination zone, address, or user.
	• Operator —Specify the operator to relate the attribute to a value.
	• Value—Specify the value to match.
Negate	Select the check box to apply the negation of the specified query, meaning that the report will contain all information that is not a result of the defined query.

Table 123. Botnet Report Settings

Managing PDF Summary Reports

▶ Monitor > PDF Reports > Manage PDF Summary

PDF summary reports contain information compiled from existing reports, based on data for the top 5 in each category (instead of top 50). They also contain trend charts that are not available in other reports.



Figure 31. PDF Summary Report

To create PDF summary reports, click **Add**. The **Manage PDF Summary Reports** page opens to show all of the available report elements.

PDF Summary Report		0
Name Summary Report 1		
🕅 Application Reports 👻 🕅 Threat Reports 👻 🕅	Traffic Reports 👻 间 Trend Reports 👻 间 URL Filtering	g Reports 👻 <u> </u> Custom Reports 👻
Bandwidth trend (Bar Graph) 🗙	Top denied sources	Top security rules
Risk trend (Line Graph)	Top destination countries	Top source countries
Threat trend (Bar Graph)	Top destination zones	Top source zones
Top connections	Top destinations	Top sources
Top denied applications	Top egress interfaces	Top unknown TCP connections 🔀
Top denied destinations	Top ingress interfaces	Top unknown UDP connections 🗙
		OK Cancel

Figure 32. Managing PDF Reports

Use one or more of these options to design the report:

- To remove an element from the report, click the \Box icon in the upper-right corner of the element's icon box or remove the check box from the item in the appropriate drop-down list box near the top of the page.
- Select additional elements by choosing from the drop-down list boxes near the top of the page.
- Drag and drop an element's icon box to move it to another area of the report.



Note: A maximum of 18 report elements is permitted. You may need to delete existing elements to add additional ones.

Click **Save**, enter a name for the report, as prompted, and click **OK**.

To display PDF reports, choose **PDF Summary Report**, and select a report type from the dropdown list at the bottom of the page to display the generated reports of that type. Click an underlined report link to open or save the report.

Managing User Activity Reports

► Monitor > PDF Reports > User Activity Report

Use this page to create reports that summarize the activity of individual users. Click **New** and specify the following information.

Description
Enter a name to identify the report (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Enter the user name or IP address (IPv4 or IPv6) of the user who will be the subject of the report.
Select the time frame for the report from the drop-down list.

Table 124. User Activity Report Settings

To run the report on demand, select the report and click **Edit**, and then click **Run**.

Managing Report Groups

Monitor > PDF Reports > Report Groups

Report groups allow you to create sets of reports that the system can compile and send as a single aggregate PDF report with an optional title page and all the constituent reports included.

Field	Description
Name	Enter a name to identify the report group (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Title Page	Select the check box to include a title page in the report.
Title	Enter the name that will appear as the report title.
Report selection	Select reports from the left column and click Add to move each report to the report group on the right. You can select Predefined, Custom, PDF Summary, and Log View report types.
	The Log View report is a report type that is automatically created each time you create a custom report and uses the same name as the custom report. This report will show the logs that were used to build the contents of the custom report.
	To include the log view data, when creating a report group, you add your custom report under the Custom Reports list and then add the log view report by selecting the matching report name from the Log View list. When you receive the report, you will see your custom report data followed by the log data that was used to create the custom report.

Table 125. Report Group Settings

To use the report group, refer to "Scheduling Reports for Email Delivery" in the next section.

Scheduling Reports for Email Delivery

Monitor > PDF Reports > Email Scheduler

Use the Email scheduler to schedule reports for delivery by email. Before adding a schedule, you must define report groups and an email profile. Refer to "Managing Report Groups" on page 272 and "Configuring Email Notification Settings" on page 83.

Scheduled reports begin running at 2:00 AM, and email forwarding occurs after all scheduled reports have finished running.

Field	Description
Name	Enter a name to identify the schedule (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Report Group	Select the report group (refer to "Managing Report Groups" on page 272).
Recurrence	Select the frequency at which to generate and send the report.
Email Profile	Select the profile that defines the email settings. Refer to "You can configure a custom log format in a Syslog Server Profile by selecting the Custom Log Format tab in Device > Server Profiles > Syslog. Click the desired log type (Config, System, Threat, Traffic, or HIP Match) and then click the fields you want to see in the logs. The tables that follow shows the meaning of each field for each log type." on page 77 for information on defining email profiles.
Override Recipient email(s)	Enter an optional email address to use instead of the recipient specified in the email profile.

Table 126. Email Scheduler Settings

Viewing Reports

Monitor > Reports

The firewall provides various "top 50" reports of the traffic statistics for the previous day or a selected day in the previous week.

To view the reports, click the report names on the right side of the page (Custom Reports, Application Reports, Traffic Reports, Threat Reports, URL Filtering Reports, and PDF Summary Reports).

By default, all reports are displayed for the previous calendar day. To view reports for any of the previous days, select a report generation date from the **Select** drop-down list at the bottom of the page.

The reports are listed in sections. You can view the information in each report for the selected time period. To export the log in CSV format, click **Export to CSV**. To open the log information in PDF format, click **Export to PDF**. The PDF file opens in a new window. Click the icons at the top of the window to print or save the file.

Generating Custom Reports

Monitor > Manage Custom Reports

You can create custom reports that are optionally based on existing report templates. The reports can be run on demand or scheduled to run each night. To view previously defined reports, choose **Reports** on the side menu.

Click **Add** to create a new custom report. To base a report on an existing template, click **Load Template** and choose the template.

Specify the following settings to define the report.

Field	Description
Name	Enter a name to identify the report (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Database	Choose the database to use as the data source for the report.
Time Frame	Choose a fixed time frame or choose Custom and specify a date and time range.
Sort By	Choose sorting options to organize the report, including the amount of information to include in the report. The available options depend on the choice of database.
Group By	Choose grouping options to organize the report, including the amount of information to include in the report. The available options depend on the choice of database.
Scheduled	Select the check box to run the report each night. The report then becomes available by choosing Reports on the side menu.
Columns	Choose columns to include in the report from the Available column and use the right-facing arrows to move them to the Selected column. Use the up and down arrows to reorder the selected columns, and use the left-facing arrows to remove previously selected columns.

Table 127. Custom Report Settings

Field	Description				
Query and Query Builder	To build a report query, specify the following and click Add . Repeat as needed to construct the full query.				
	• Connector —Choose the connector (and/or) to precede the expression you are adding.				
	• Attribute—Choose a data element. The available options depend on the choice of database.				
	• Operator —Choose the criterion to determine whether the attribute applies (such as =). The available options depend on the choice of database.				
	• Value —Specify the attribute value to match.				
	For example, the following figure (based on the Traffic Log database) shows a query that matches if the traffic log entry was received in the past 24 hours and is from the "untrust" zone.				
	Query (receive_time in last-24-hrs) and (zone eq untrust)				
	Query Builder Connector Attribute Operator Value and Zone = untrust or Zone.src # Negate zone.dst				
Negate	Select the check box to interpret the query as a negation. In the previous example, the negate option causes a match on entries that are not in the past 24 hours or are not from the "untrust" zone.				

 Table 127.
 Custom Report Settings (Continued)

Identifying Unknown Applications and Taking Action

There are several ways to view unknown applications using the web interface of the Palo Alto Networks devices:

- **Application Command Center (ACC)**—Unknown applications are sorted along with other applications in the ACC. Click a link for an unknown application to view the details of the application, including top sources and destinations. For top sources, click the
 - link to look up the owner of the address.

							Ap	plications
	Risk	Application	n Name	Sessi	ons	Ву	tes	Thr
1	1	insufficien	t-data	36.0 K		2.6 M		0
2	2	dns		9.3K 🛄		2.4 M		0
3	4	ssl		7.1K 🛄		279.5 M		0
4	4	web-brow	sing	4.6 K 🔲		131.3 M		37 🔳
5	1	panos-we	b-interface	627 📘		14.6 M		0
6	3	ms-ds-smb	5	567		461.3 M		433
7	1	unknown-	tcp	29		5.6 M		0
8	2	kerberos		499		1.8 M		0
9	2	ntp		359 🛿		58.4K		o 🛙
1	10, 16, 0, 4	13	10.16.0.43	paloaltonetwor	rk∖jfitz 4.6 M 🗖	bytes	10	33013
1	10.16.0.4	3	10. 16.0.43 📽	paloaltonetwo	*k\jfitz 4.6 M		10	
2	10.20.0.4	7	pan00279.paloaltonetwork	s.local 🖓	499.1K 🔲		1 🗖	
	Destinatio	n address	Destination Host Name	Destination Us	er	Bytes	Se	essions
1	64.30.236	5.36	lps-vip1.tm.cbsig.net 🚱		229.7K 🛙		5	
2	64.30.23	5.228	ws1620-fe.tm.cbsig.net		7.7K 🛽		2	
3	68, 13, 75,	224	ip68-13-75-224.om.om.co	c.net 🚰	4.2 M		1	
4	10.16.0.2	11	mdm.paloaltonetworks.com		499.1 K 🔲		1	
5	155.41.16	51.53	dot1x-155-41-161-53.bum	c.bu.edu 🕼	161.0 K		1	
6	64.30.235	5.227	ws1619-fe.tm.cbsig.net 😫	· • •	3.8 K		1	
_	ource Cour	ntries						
	ource cour				120.00		S	
op S	Source Co	untry			Bytes			essions
op S 1	Source Co Unknown	untry		4.6 M E	Bytes		10	essions
op S 1 2	Source Co Unknown 10.0.(untry 	255.255	4.6 M E 499.1 K E	Bytes		10	
op S 1 2 0p D	Source Co Unknown 10.0.(ountry 0.0-10.255. Countries	255.255	4.6 M E 499.1 K E	Bytes		10	
op S 1 2 0p D	Source Co Unknown 10;0.(Destination	ountry 0.0-10.255. Countries n Country	255.255	4.6 M E 499.1K E	Bytes	_	10 1	essions
эр S 1 2 эр D 1	Source Co Unknown 10,0.0	ountry 0.0-10.255. Countries n Country d States	255.255	4.6 M E 499.1 K E 400.1 K E 4.6 M E	Bytes		10 1 S 10	essions

Link to look up owner of the address

Figure 33. Unknown Applications in the ACC List

• Unknown application reports—Unknown application reports are automatically run on a daily basis and stored in the Reports section of the **Monitor** tab. These reports can provide useful information to help identify unknown applications.

• **Detailed traffic logs**—You can use the detailed traffic logs to track down unknown applications. If logging is enabled for the start and end of session, the traffic log will provide specific information about the start and end of an unknown session. Use the filter option to restrict the display to entries that match "unknown-tcp," as shown in the next figure.

٩ (app eq unknown-tc;)								🖃 🗶 🕂	
	Receive Time	Туре	From Zone	To Zone	Source	Source User	Destination	To Port	Application	Action	Rule
Þ	09/12 19:42:52	end	trust	untrust	10.16.0.43	paloaltonetwork\jfitz-gerald	64.30.236.37	2112	unknown-tcp	allow	Let PM do
Þ	09/12 18:52:11	end	trust	untrust	10.16.0.203	paloaltonetwork\savasarala	98.137.88.84	80	unknown-tcp	allow	Let PM do
D I	09/12 18:51:50	end	trust	untrust	10, 16, 0, 203	paloaltonetwork\savasarala	209.73.190.208	80	unknown-tcp	allow	Let PM do
p ;	09/12 18:51:50	end	trust	untrust	10.16.0.203	paloaltonetwork\savasarala	209.73.190.208	80	unknown-tcp	allow	Let PM do
Þ	09/12 18:51:50	end	trust	untrust	10, 16, 0, 203	paloaltonetwork\savasarala	209.73.190.208	80	unknown-tcp	allow	Let PM do
Ø	09/12 18:51:50	end	trust	untrust	10.16.0.203	paloaltonetwork\savasarala	209.73.190.208	80	unknown-tcp	allow	Let PM do
Þ	09/12 18:38:21	end	trust	untrust	10, 16, 0, 203	paloaltonetwork\savasarala	209.73.190.208	80	unknown-tcp	allow	Let PM do
ð	09/12 18:38:21	end	trust	untrust	10.16.0.203	paloaltonetwork\savasarala	209.73.190.208	80	unknown-tcp	allow	Let PM do
ð	09/12 18:38:21	end	trust	untrust	10.16.0.203	paloaltonetwork\savasarala	209.73.190.208	80	unknown-tcp	allow	Let PM do

Figure 34. Unknown Applications in Traffic Logs

Taking Action

You can take the following actions to deal with unknown applications:

- Use custom application definition with application override (refer to "Custom Application Definition with Application Override" on page 205).
- Use custom applications with signatures (refer to "Custom Applications with Signatures" on page 236).
- Request an App-ID from Palo Alto Networks (refer to "Requesting an App-ID from Palo Alto Networks" in the next section).

Policies can also be set to control unknown applications by unknown TCP, unknown UDP or by a combination of source zone, destination zone, and IP addresses. Refer to "Application Override Policies" on page 205.



Note: You can use custom signatures in App-ID definitions.

Requesting an App-ID from Palo Alto Networks

If it is necessary to identify an application using application contents instead of port, protocol, and IP address, you can submit the application to Palo Alto Networks for classification. This is important for applications that run over the Internet and for which custom application does not work. You can submit the application to Palo Alto Networks in either of the following ways:

• If the application is a readily accessible on the Internet (for example, an instant messaging application), then submit the name of the application and the URL to your account team or to this web site: *http://www.paloaltonetworks.com/researchcenter/tools*.

• If the application is not easily accessible (for example, a customer relationship management application) you must submit a packet capture (PCAP) of the running application using the session packet capture function built into the firewall. For assistance, contact technical support at *support@paloaltonetworks.com*.

Other Unknown Traffic

The firewall may report an application to be "unknown" in the ACC, logs, or reports for either of the following reasons:

- **Incomplete**—A handshake took place, but no data packets were sent prior to the timeout.
- **Insufficient-Data**—A handshake took place followed by one or more data packets; however, not enough data packets were exchanged to identify the application.

Taking Packet Captures

Monitor > Packet Capture

PAN-OS supports packet capture for troubleshooting or detecting unknown applications. You can define filters such that only the packets that match the filters are captured. The packet captures are locally stored on the device and are available for download to your local computer.



Note: Packet Capture is for troubleshooting only. This feature can cause the system performance to degrade and should be used only when necessary. After the capture is complete, please remember to disable the feature.

To specify filtering and capture options, specify the information in the following table. To clear all filtering and capture settings, click **Clear All Settings**.

To select capture files for download, click the file name in the capture file list on the right side of the page.

Field	Description
Filtering	
Manage Filters	Click Manage Filters , click Add to add a new filter, and specify the following information:
	• Id—Enter or select an identifier for the filter.
	• Ingress Interface—Select the firewall interface.
	• Source —Specify the source IP address.
	• Destination—Specify the destination IP address.
	• Src Port —Specify the source port.
	• Dest Port—Specify the destination port.
	• Proto —Specify the protocol to filter.
	• Non-IP—Choose how to treat non-IP traffic (exclude all IP traffic, include all IP traffic, include only IP traffic, or do not include an IP filter).
	• IPv6—Select the check box to include IPv6 packets in the filter.

Table 128. Packet Capture Settings

Field	Description
Filtering	Click to toggle the filtering selections on or off.
Pre-Parse Match	Click to toggle the pre-parse match option on or off.
	The pre-parse-match option is added for advanced troubleshooting purposes. After a packet enters the ingress port, it proceeds through several processing steps before it is parsed for matches against pre- configured filters.
	It is possible for a packet, due to a failure, to not reach the filtering stage. This can occur, for example, if a route lookup fails.
	Set the pre-parse-match setting to ON to emulate a positive match for every packet entering the system. This allows the firewall to capture even the packets that do not reach the filtering process. If a packet is able to reach the filtering stage, it is then processed according to the filter configuration and discarded if it fails to meet filtering criteria.
Capture Files	
Capturing	Click to toggle packet capturing on or off.
Capture Settings	Click Add and specify the following:
	 Stage—Indicate the point at which to capture the packet:
	 drop—When packet processing encounters an error and the packet is to be dropped.
	 firewall—When the packet has a session match or a first packet with a session is successfully created.
	- receive —When the packet is received on the dataplane processor.
	 transmit—When the packet is to be transmitted on the dataplane processor.
	• File—Specify the capture file name. The file name should begin with a letter and can include letters, digits, periods, underscores, or hyphens.
	 Packet Count—Specify the number of packets after which capturing stops.
	• Byte Count —Specify the number of bytes after which capturing stops.

Table 128. Packet Capture Settings (Continued)

Taking Packet Captures

Chapter 7 Configuring the Firewall for User Identification

This chapter describes how to configure the firewall to identify the users who attempt to access the network.

- "Overview of User Identification" in the next section
- "User Identification Agents" on page 284
- "PAN-OS User Mapping Configuration" on page 291
- "Setting Up the User-ID Agent" on page 296
- "Setting Up the Terminal Services Agent" on page 301

Overview of User Identification

User Identification (User-ID) is a feature of Palo Alto Networks firewalls that allows administrators to configure and enforce firewall policies based on users and user groups, instead of or in addition to network zones and addresses.

User-ID identifies the user on the network and the IP addresses of the computers the user is logged into to effectively enforce firewall policies. User-ID can also retrieve user and group information from a connected LDAP directory, allowing administrators to configure policies based on user groups, which are then translated into a list of users.

How User Identification Works

The functionality provided by User-ID requires the collection of information from the network and directory servers. The following elements are involved in the information collection:

• Identifying users on the network

User-ID provides a variety of mechanisms to reliably identify network users and their associated login session information (computers and network addresses). Some of the mechanisms require the installation of a User-ID Agent on network servers to provide the

most transparent user experience. With PAN-OS 5.0 and greater, you can use the built in PAN-OS User Mapping feature that performs the same function as the User-ID Agent, but does not require an agent installation.

• Event log monitoring

Whenever a user authenticates to the Active Directory (AD) domain, a Microsoft Windows server, or Microsoft Exchange server, an event log is produced. Users can be identified on the network by monitoring those servers for the corresponding login events.

• Server session monitoring

Another method is to continually monitor servers for network sessions established by users on the network. When a user successfully authenticated to a server, the session table of the server provides the user name and network source the user is connecting from.

• Client Probing

In a Microsoft Windows environment, the client system can provide information about logged on users through Windows Management Instrumentation (WMI) for authorized users and services. Probing Microsoft Windows clients on demand provides information on users logged into a client computer.

• XML API

Other identification methods are not directly supported by the User-ID features and options. For these cases, an XML over SSL interface is available, allowing customized solutions to register valid users and their corresponding client address on the network with User-ID.

• Captive Portal

If the user cannot be identified based on login information, an established session or client probe, the firewall can re-direct any outbound HTTP requests and re-direct the user to a web form. The web form can transparently authenticate the user through a NTLM challenge, which is automatically evaluated and answered by the web browser or through an explicit login page.

• Shared computers

Shared computers, such as Microsoft Terminal Servers, are problematic for most implementations, because a number of users share the same system and therefore the same network address. In this case, an Agent can be installed on the Terminal Server, which then associates not just the network address, but also allocated port ranges to the logged in users.

Identifying Users and Groups

Policy management on the basis of individual users is unmanageable; therefore, users need to be associated and tied to user groups. Every enterprise environment stores user information in a directory service, such as Microsoft Active Directory or Novell eDirectory. All of those directory services are accessible through LDAP or LDAP over SSL (LDAPS).

The directory services provide resolution of user names and the associated user groups, which allows firewall administrators to configure security policies for user groups rather than individual users.

How User-ID Components Interact

The User-ID Agent, PAN-OS, and the Terminal Services Agent interact with each other to provide complete user identification services.

User-ID Agent

The User-ID Agents identifies the user on the network using one or all of the mechanisms described previously in this chapter.

• Gathering user and login information

The User-ID Agent can be configured to monitor up to 100 Microsoft Windows Servers for user login events. When the Agent first connects to a server, it automatically retrieves a list of the last login events from the domain controller. During normal operations, it continues to receive new event information. The User-ID Agent provides the collected information to the firewall to enforce policy based on users and groups.

• Providing users and network addresses to connected devices

To provide user and network address information, the firewall establishes a persistent connection to the User-ID Agent and retrieves a list of all identified users and network addresses on its first connection and every hour. During each hour, the firewall retrieves changes that the Agent detects.

• On demand user identification

If the firewall identifies a new network address in the network traffic for which no user is listed, it can contact the User-ID Agent and request it to identify the user. This is done through a WMI or NetBIOS probe to the specific network address. When the client identifies the user, a new network address and username association is created and provided to the firewall.

PAN-OS User Mapping

The PAN-OS User Mapping feature performs the same functions as the User-ID Agent, but does not require an agent installation on domain servers. This feature communicates directly with Microsoft servers and Novell eDirectory servers to gather user to IP and user to group mapping information. The enumeration of the individual users in a user group is also performed on the firewall. LDAP search and NTLM authentication is also supported. Refer to "User Mapping Tab" on page 286 and "PAN-OS User Mapping Configuration" on page 291.

You can also configure the firewall to share the user mapping information with other PAN-OS firewalls in your network. This is useful if you want to designate one or more devices to perform user mapping functions and then share that information with other firewalls. By using the include and exclude lists, you can also filter certain networks from the collection process. Refer to "PAN-OS User Mapping Configuration" on page 291.

Terminal Services Agent

The Terminal Services Agent (TS Agent) solves the problem of multiple users using the same machine at the same time, for example on a Microsoft Terminal Server. After it is installed on the server, it allocates specific port ranges to each individual user. Every user connection is established using a port within the specific allocated port range.

When a port range is allocated for a particular user, the Terminal Services Agent notifies every connected firewall about the allocated port range so that policy can be enforced based on user and user groups.

PAN-OS LDAP Group Query

In addition to enforcing policy based on individual users, the firewall can also be configured to allow or block traffic for groups of users. The enumeration of the individual users in a user group is performed by the firewall.

For this purpose, a LDAP server entry and group mapping settings need to be configured. The resulting LDAP query retrieves user groups and the corresponding list of group members.

This operation is performed every time a new configuration is submitted. Changes in group membership are detected through specific LDAP searches that retrieve only the groups and their member list that changed since the last search was performed.

User Identification Agents

The User Identification Agent (User-ID Agent) is a Palo Alto Networks application that is installed on your network to obtain needed mapping information between IP addresses and network users. The User-ID Agent collects user-to-IP address mapping information from the domain controller security logs and provides it to the firewall for use in security policies and logs.

Refer to "PAN-OS User Mapping" on page 283, which performs the same functions as the User-ID Agent, but does not require an agent to be installed on the domain servers.

The IP address-to-user name mapping relies on the following mechanisms:

- For Active Directory, the security logs are continually monitored on the domain controller to detect user login events that contain user and IP address information.
- For Active Directory, a direct connection is required to all Domain Controllers to monitor user session activity and determine the user IP addresses.
- For eDirectory, when a user logs in, the IP address information is stored in eDirectory and retrieved by the User-ID Agent.
- The host PC is polled to verify IP address and user information using Windows Management Instrumentation (WMI) or Network Basic Input/Output System (NetBIOS). This occurs every 20 minutes to verify that the IP address-to-user name mapping is still correct and also when an IP address is seen that does not have an associated user name.
- The User-ID Agent application programming interface (API) is used to send information on user IP addresses to the User-ID Agent.
- User group mapping is performed through LDAP queries on directory servers. The firewall performs LDAP queries directly, but can use a configured User-ID agent as a LDAP proxy in cases where caching is desirable or direct access from the firewall to the directory server is not possible."



Note: User identification mapping requires that the firewall obtain the source IP address of the user before the IP address is translated with NAT. If multiple users appear to have the same source address, due to NAT or use of a proxy device, accurate user identification is not possible.

In addition to the User-ID Agents, the firewall supports a Terminal Services Agent (TS Agent) that allows the firewall to identify individual users who are supported by the same terminal server. The firewall also supports captive portals for situations in which the User-ID Agent is unable to associate a user with an IP address.

Refer to the following sections for further information:

- "Captive Portals" in the next section
- "Configuring the Firewall for User Identification" on page 286
- "Setting Up the User-ID Agent" on page 296
- "Setting Up the Terminal Services Agent" on page 301

Captive Portals

If the User-ID Agent is unable to associate a user with an IP address, a captive portal can take over and authenticate the user with a web form or NT LAN Manager (NTLM) challenge.

To receive the web form, users must be using a web browser and be in the process of connecting. Upon successful authentication, users are automatically directed to the originally requested web site. The firewall can now execute policies based on the user information for any applications passing through the firewall, not just for applications that use a web browser.

The following rules apply to captive portals:

- Captive portal rules work for HTTP and HTTPS web traffic.
- If the action for the rule is "web form," a web form is presented to the user to prompt for a password.
- If the rule is "NTLM" and the browser is Internet Explorer or Firefox, the firewall performs an NTLM authentication challenge (transparent to the user). If another browser is used, the web form is presented.

If the above-mentioned captive portal rules do not apply because the traffic is not HTTP or there is no rule match, then the firewall applies its IP-based security policies (as opposed to user-based security policies).

Configuring the Firewall for User Identification

Device > User Identification

Use the settings on this page to configure the firewall for user identification.

- User Mapping tab—Specify settings to use the PAN-OS User Mapping feature to provides accurate mappings between IP addresses and logged in users as well as user to group membership mapping. This option performs the same functions as the User-ID Agent directly from the firewall, so no agent is required on the domain controllers.
- User-ID Agents tab—Specify settings to support the user identification agent, which provides accurate mappings between IP addresses and logged in users.
- **Terminal Services tab**—Specify settings to support the terminal services agent. Refer to "Setting Up the Terminal Services Agent" on page 301.
- **Group Mappings Settings tab**—Specify settings to support mappings that associate users with user groups. User group mapping is performed by the firewall.
- **Captive Portal tab**—Specify settings to support use of a captive portal for user identification. Refer to "Captive Portals" on page 285.

Field	Description
User Mapping Tab	
WMI Authentication tab	User Name —Specify the domain\username account that has permissions to perform WMI queries on client computers.
	Password/Confirm Password—Specify the account password.
Server Monitor tab	Enable Security Log —Select the check box to enable security log monitoring on Windows servers. Security logs will be queried to locate user to IP mapping information on the servers specified in the Server Monitoring list.
	Server Log Monitor Frequency (sec) —Specify how often the firewall will query Windows servers for user to IP mapping information (default 1 second, range 1-3600 seconds).
	Enable Session —Select the check box to enable monitoring of user sessions on the servers specified in the Server Monitoring list. Each time a user connects to a server, a session is created and this information can also be used to identify the user IP address.
	Server Session Read Frequency (sec) —Specify how often the firewall will query Windows server user sessions for user to IP mapping information (default 10 second, range 1-3600 seconds).
	Novell eDirectory Query Interval (sec) —Specify how often the firewall will query Novell eDirectory servers for user to IP mapping information (default 30 second, range 1-3600 seconds).

Table 129. User-ID Agent Settings

Field	Description
Client Probing tab	Enable Probing —Select this check box to enable WMI/NetBIOS probing to each client PC identified by the user mapping process. Probing will help ensure that the same user is still logged into the client PC in order to provide accurate user to IP information.
	Probe Interval (min) —Specify the client PC probe interval (default 20 minutes, range 1-1440 minutes).
	In large deployments, it is important to set the probe interval properly to allow time to probe each client that has been identified. Example, if you have 6,000 users and an interval of 10 minutes, it would require 10 WMI request a second from each client.
	Note: For WMI polling to work effectively, the User Mapping profile must be configured with a domain administrator account, and each probed client PC must have a remote administration exception configured in the Windows firewall. For NetBIOS probing to work effectively, each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled.
Cache tab	Enable User Identification timeout —Select this check box to enable a timeout value for user to IP mapping entries. When the timeout value is reached, the user to IP mapping will be cleared and a new mapping will be collected. This will ensure that the firewall has the most current information as users roam around and obtain new IP addresses.
	User Identification Timeout (min) —Set the timeout value for user to IP mapping entries (default 45 minutes, range 1440).
NTLM tab	 Enable NTLM authentication processing—Select this check box to enable NT LAN Manager (NTLM) authentication processing. When Captive Portal is set up with a web form to capture user to IP information, the client will be transparently authenticated through an NTLM challenge. With this option enabled, the firewall will collect this information from the NTLM domain. If both the PAN-OS User-ID collector and the User-ID Agent is installed on domain controllers, NTLM responses will go directly to the domain controller. When the firewall is configured to share its User-ID information with other firewalls, it can serve NTLM requests coming from other PAN-OS firewalls, performing the function of the User-ID
	Agent.
	Admin User Name—Enter the administrator account that has access to
	Password/Confirm Password —Enter the password for the administrator
	Note: You cannot enable this check box in more than one virtual system.

Table 129. User-ID Agent Settings (Continued)

Field	Description
Redistribution tab	Collector Name —Specify the collector name if you want this firewall to act as a user mapping redistribution point for other firewalls on your network.
	The collector name and pre-shared key are used when configuring the User-ID Agents on the firewalls that will pull the user mapping information.
	To enable a firewall to act as a re-distribution point, you also need to enable the
	User-ID Service in Network > Network Profiles > Interface Mgmt .
	Pre-Shared Key/Confirm Pre-Shared Key —Enter the pre-shared key that is used by other firewalls to establishing a secure connection for user mapping transfers.
Server Monitoring	Click Add to configure the domain server(s) that the firewall will contact to gather user IP mapping and user to group mapping information.
	Name—Enter the domain server network address.
	Network Address—Enter the network address of the domain server.
	Type —Select the type of domain server (Microsoft Active Directory, Microsoft Exchange, or Novell eDirectory).
	Enable —Select the check box to enable the domain server profile.
	Discover —Click this option to discover Microsoft Active Directory servers via DNS. You can then select check box next to the servers you want to use to obtain user mapping information. The firewall will discover domain controllers based on the domain name entered in the Device > Setup > Management tab > General Settings page Domain field.
Include/Exclude Networks	Name —Enter a name to identify the profile that will include or exclude a network for User-ID discovery purposes. This option allows you to include or exclude a network range for IP address-to-user name mapping. Example, if you exclude 10.1.1.0/24, User-ID will not try to find user names for IP addresses in the excluded range. This in turn will also include or exclude ranges for mappings sent to other PAN-OS firewalls.
	When defining an include or exclude network range, an implicit exclude- all will be performed. For example, if you include 10.1.1.0/24, all other networks will be excluded. If you exclude 10.1.1.0/24, all networks will be excluded, so when using exclude you must also have an include network, otherwise all networks are excluded.
	Enabled —Select this option to enable the include/exclude profile.
	Discovery —Select the option to either Include or Exclude the defined network range.
	Network Address —Enter a network range that you would like to include or exclude for IP address-to-user name mapping discovery. For example, 10.1.1.0/24.
User-ID Agents Tab	
Name	Enter a name to identify the User-ID Agent (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.

Table 129. User-ID Agent Settings (Continued)
Field	Description
Virtual System	Select the virtual system from the drop-down list (if supported on the firewall model).
IP Address	Enter the IP address of the Windows PC on which the User-ID Agent is installed.
Port	Enter the port number on which the User-ID Agent service is configured on the remote host.
	When using a PAN-OS firewall as a redistribution point for user mapping information and you configure the User-ID Agent to connect to that device, you must use port 5007.
Collector Name	Enter the collector name that the User-ID Agent will pull from to obtain user mapping information. Refer to the "User Mapping Tab" on page 286.
Collector Pre-shared Key/Confirm Collector Pre-shared key	Enter the pre-shared key that will be used to allow SSL connectivity between the User-ID Agent and the firewall that is acting as a distribution point for user mapping.
Use as LDAP Proxy	Select the check box if the User-ID Agent is to be used as a LDAP proxy instead of the firewall connecting directly to the directory service.
Use for NTLM Authentication	Select the check box to use the configured User-ID Agent to verify NTLM client authentication from the captive portal with the Active Directory domain.
Enabled	Select the check box to enable the user identification agent.
Refresh Connected	Click to refresh the connection status of User-ID Agent profiles. This button is located to the right of the Add and Delete buttons.
Custom Agent Sequence	This option allows you to define the sequence order in which the User-ID agent profiles will connect to the defined server. For example, if you have four agents identified in the sequence list, it will attempt to connect to the first agent listed, if that connection fails, it will connect to the next agent listed, and so on. If this option is not configured, the connection sequence will follow the order of the agents listed in the main page.
Terminal Services Agent Tab	
Name	Enter a name to identify the TS Agent (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Virtual system	Select the virtual system from the drop-down list (if supported on the firewall model).
Host	Enter the IP address of the Windows PC on which the TS Agent will be installed. You can also specify alternative IP addresses (see the last entry in this table).
Port	Enter the port number on which the User-ID Agent service is configured on the remote host.
Alternative IP Addresses	Enter additional IP addresses, if the server has multiple IP addresses that can appear as the source IP address for the outgoing traffic.

Table 129. User-ID Agent Settings (Continued)

Field	Description
Group Mapping Settings Tab	
Name	Enter a name to identify the user-to-group mapping for user identification (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Virtual system	Select the virtual system from the drop-down list (if supported on the firewall model).
Server Profile subtab	Specify the following settings:
	• Select an LDAP server profile from the drop-down list, and specify the interval (seconds) after which the configuration is updated with any new server profile information.
	• Group objects
	 Search Filter—Specify an LDAP query that can be used to control which groups are retrieved and tracked.
	 Object Class—Specify the definition of a group. For example, the default is objectClass=group, which means that the system retrieves all objects in the directory that match the group filter and have object- Class=group.
	 Group Name—Enter the attribute that specifies the name of the group. For example in Active Directory, this attribute is "CN" (Common Name).
	 Group Member—Specify the attribute that contains the members of this group. For example in Active Directory, this attribute is "member."
	• User Objects
	 Search Filter—Specify an LDAP query that can be used to control which users are retrieved and tracked.
	 Object Class—Specify the definition of the a user object. For example in Active Directory, the objectClass is "user."
	 User Name—Specify the attribute for user name. For example in Active Directory, this attribute is "samAccountName."
Group Include List	Locate groups in the Available Groups list. Click the $\textcircled{1}$ icon to add the groups to the Included list and click the $\textcircled{1}$ icon to remove groups from the list.
Captive Portal Settings Tab	
Enable Captive Portal	Select to enable the captive portal option for authentication.
Location	Select the virtual system from the drop-down list (if supported on the firewall model).
Idle Timer	Enter the length of time after which the captive portal page times out (1-1440 minutes, default 15 minutes).
Expiration	Specify the timeout interval (range 1 - 1440 minutes, default 60 minutes).
Redirect Host	Specify the hostname used for the HTTP redirect used to initiate the NTLM challenge sent to the client.

Table 129. User-ID Agent Settings (Continued)

Field	Description
Server Certificate	Select the HTTP SSL certificate used for captive portal. Note: If you select None , the firewall will use the local default certificate to provide an SSL connection.
Certificate	Choose the certificate profile to use for client authentication.
Authentication Profile	Choose the profile to determine the authentication source for captive portal logins.
NTLM Authentication	For NTLM authentication, specify the following:
	 Attempts—Specify the number of attempts after which the NTLM authentication fails.
	• Timeout —Specify the number of seconds after which the NTLM authentication times out.
	• Reversion Time —Specify the time after which the firewall will again try to contact the first agent in the list of User-ID Agents after the agent becomes unavailable.
	Note: These options only apply to the User-ID Agent installed on domain servers and does not apply to the PAN-OS User-ID feature configured in the User Mapping tab on the User Identification page.
Mode	Choose whether the captive portal will use redirection or be transparent to the user.
	Redirection is required for NTLM and session cookie retention. With the redirection option, the firewall can set a cookie for future login requests. Future redirection then becomes transparent to the user if the browser has not been closed.
	For session cookies, specify the following settings:
	• Enable —Select the check box to configure an interval after which the redirection times out.
	• Timeout —If Enable is selected, specify the timeout interval (range 60 - 10080 minutes, default 1440 minutes).
	• Roaming —Select the check box if to retain the cookie if the IP address changes while the browser is open (for example, if the client moves from a wired to wireless network). The cookie is lost when the browser closes, whether or not Roaming is selected.
	Note: To use the captive portal in redirect mode, you must enable response pages on the interface management profile assigned to the Layer 3 interface to which you are redirecting the active portal. Refer to "Defining Interface Management Profiles" on page 175 and "Configuring Layer 3 Interfaces" on page 130.

Table 129. User-ID Agent Settings (Continued)

PAN-OS User Mapping Configuration

This section describes the basic steps needed to configure a firewall to retrieve IP address-touser name mapping data directly from domain servers. This feature does not require the installation of a User-ID Agent on the domain servers. The firewall can also be configured to redistribute the user mapping information to other firewalls.

Refer to the following sections:

• "Configuring PAN-OS User Mapping" in the next section

• "Configure a Firewall to Share User Mapping Data" on page 295

Configuring PAN-OS User Mapping

This section describes the configuration information needed to gather IP address-to-user name mapping and group mapping data for User Identification.

In this example, there is no need to install a User-ID Agent client on the domain controller. The following components need to be configured:

- **Domain Account**—A domain account is needed to query Microsoft Active Domain Controllers to obtain user information from security logs and to perform LDAP queries.
- User-ID Agent—As of PAN-OS 5.0, you configure the User-ID agent using the User Mapping tab, which is used to configure the native version of the agent. In large environments where performing user mapping enumeration may start to cause performance issues on the firewall, you will want to use the agent installed on domain controllers, which is configured in the User-ID Agents tab. For information on installing this agent, refer to "Setting Up the User-ID Agent" on page 296.
- **LDAP Server Profile**—This profile is used to define the domain controller server information that the firewall will use to perform queries for gathering user and group mapping information.
- **Group Mappings**—These settings allow the firewall to associate users with user groups, which enables you to create policies based on directory groups. The LDAP Server Profile is used in the Group Mappings settings to define the servers that will be queried.
- **Policies**—You need to configure the appropriate security profiles so the firewall can communicate with the defined directory servers.

The following outlines the basic steps needed to configure PAN-OS User Mapping:

1. Obtain a login account and password that has permissions to the servers and hosts that you will use to retrieve user mapping information. The permissions needed is based on the methods you will enable to gather the user mapping information, which may be from

Microsoft Active Directory servers, Microsoft Exchange servers, Microsoft file servers, Novell eDirectory, or using WMI to access client PCs.



Note: In Windows 2008 or later domains, you can add an account to the "Event Log Readers" group to access event logs to obtain IP address-touser name mapping information from event logs. In Windows 2003 domains, the account should be assigned "Manage Auditing and Security logs" permissions through group policy.

You can also add the account to Server Operators to grant permissions to view current open server sessions to obtain IP address-to-user name mapping information for Windows 2003 and 2008 domains.

For the LDAP group queries to identify user to group mappings, no special membership is needed, most accounts can obtain group membership in Active Directory.

WMI is reliable and is secured by NTLM or Kerberos based authentication. To perform WMI queries, an account is needed that has rights to read CIMV2 namespace on the client system. By default, domain administrators have this permissions.

- 2. Navigate to **Device > User Identification** and click the **User Mapping** tab.
- 3. Click the Edit icon on the upper right side of the Palo Alto Networks User ID Agent Setup window. On this page, you define the settings that will be used when collecting IP address-to-user name mapping information on the servers defined in the Server Monitoring list. Populate each tab based on your requirements. Refer to "User Mapping Tab" on page 286 for descriptions of each of the user mapping settings. The Redistribution tab is only used if you plan on using this firewall to distribute user mapping information to other devices.
- 4. Click **OK** to save the settings.
- 5. In the Server Monitoring section, click Add to add the domain servers that you will access to gather user information. Enter the server Name, Network Address, and Type (Microsoft Active Directory, Microsoft Exchange, or Novel eDirectory). You can also click the Discover icon to discover domain controllers using DNS lookups. The firewall will discover domain controllers based on the domain name entered in the Device > Setup > Management > General Settings page.
- 6. Select the check box next to each domain server that you will monitor.
- 7. To specify networks that will be included or excluded in the user mapping data, click the Add icon in the Include/Exclude section and define the networks you want to filter. This allows you to send filtered user mapping data to the other firewalls in your network that use this device to pull user mapping information.



Note: You can also exclude a list of user accounts from being mapped by creating an ignore-user list. When using the Palo Alto Networks User Mapping feature, the list of user accounts to exclude is created in the CLI by using the command set user-id-collector ignore-user <value>.

There is no limit to the number of accounts that can be added to this list.

8. Click Commit to make the changes active.

You can check the status of the directory server connection by running the show user server-monitor state or show user server-monitor statistics commands.

Now that the User-ID Agent is configured, you need to configure an LDAP profile.

- 1. Name your configuration and leave the **Administrator use only** check box unchecked.
- 2. List the directory servers that you want the firewall to use in the server list. You need to provide at least one server; two or more are recommended for failover purposes. The standard LDAP port for this configuration is 389.
- 3. Leave the "Domain" field empty, unless you want to configure multiple independent directories.
- 4. Select a directory "Type". Based on the selected directory type, the firewall can populate default values for attributes and objectclasses used for user and group objects in the directory server.
- 5. Enter the base of the LDAP directory in the "Base" field. For example, if your Active Directory Domain is "paloaltonetworks.local", your base would be "dc=paloaltonetworks, dc=local", unless you want to leverage an Active Directory Global Catalog. In the "Bind DN" field, enter a user name for a user with sufficient permission to read the LDAP tree. In an Active Directory environment, a valid username for this entry could be the "User Principal Name", e.g. "administrator@paloaltonetworks.local" but also the users distinguished name, e.g. "cn=Administrator,cn=Users,dc=paloaltonetworks,dc=local".
- 6. Enter and confirm the authentication password for the user account that you entered above.

If you want to create policies based on directory groups, you need to configure Group Mapping. This configuration defines how groups and users are retrieved from the directory and which users groups should be included for use in policies.

- 1. Navigate to **Device > User Identification**, click the **Group Mapping Settings** tab and then click **Add**.
- 2. In the **Server Profile** drop-down, select the LDAP profile that you created in the previous section. All LDAP attributes and ObjectClasses will be pre-populated based on the directory server type you selected in the LDAP Server Profile page.
- 3. Adjust the Update Interval field if needed. The default is 3600 seconds (1 hour).
- 4. In order to optimize LDAP queries and policy configurations you can specify a list of user groups you would like to include in policies. To do this, click the **Group Include List** tab.
- 5. The left pane will show **Available Groups**, either navigate to the desired group, or search for groups that you want to include in your queries. You can use the asterisks (*) to do pattern matching in your searches.
- 6. Select the desired group(s) and then click the 🕂 symbol to add the group(s) to the **Included Groups** pane.
- 7. Click **OK** to save group mapping.

Configure a Firewall to Share User Mapping Data

Once you configure a firewall to retrieve user mapping information as described in "Configuring PAN-OS User Mapping" on page 292, you can designate the firewall to provide this information to other PAN-OS firewalls on your network.

The illustration in Figure 35 shows that a redistribution firewall communicates with the domain server to retrieve user to IP mapping information. The redistribution firewall can use any method to collect user mapping and group mapping information as described in "How User Identification Works" on page 281, and will then act as a User-ID Agent to share this information with other firewalls. Firewalls FW1, FW2, and so on, are configured to pull the mapping information directly from the redistribution firewall and do not need to communicate directly with any domain servers.



Figure 35. Firewall configured to share user mapping information

- 1. Navigate to **Device > User Identification** and click the **User Mapping** tab.
- 2. Click the **Edit** icon on the upper right side of the **Palo Alto Networks User ID Agent Setup** window and then click the **Redistribution** tab.
- 3. Enter a **Collector Name** and a **Pre-shared Key**. This information will be used by the firewall(s) that will pull user mapping information.
- 4. Click **OK** to save your changes.
- 5. Navigate to Network > Network Profiles > Interface Mgmt to enable the User-ID Agent Service.
- 6. Either **Add** a new Interface Management Profile, or modify an existing profile. Select the **User-ID Service** check box to enable the service.

The firewall is now ready to redistribute user mapping information to other firewalls.

To configure a firewall to retrieve information from a user mapping redistribution firewall, perform the following steps:

- 1. Navigate to **Device > User Identification** and click the **User-ID Agents** tab. Refer to "User-ID Agents Tab" on page 288 for setting details.
- 2. Click **Add** and enter a **Name** for the User Identification Agent profile and then populate the following fields:
 - a. **Host**—Enter the host name or IP address of the firewall that is configured as a user mapping redistribution point.

- b. **Port**—Enter 5007 for the port number. This is the designated port used for this service and cannot be modified.
- c. **Collector**—Enter the collector name of the firewall that is configured as a user mapping redistribution point.
- d. **Collector Pre-Shared Key**—Enter the collector pre-shared key of the firewall that is configured as a user mapping redistribution point.
- e. Use as LDAP Proxy—Select this check box to use the firewall as an LDAP Proxy, instead of connecting directly to the directory server. This is used to query directory group and group membership information.
- f. Use for NTLM Authentication—Select this check box to use the firewall to verify NTLM client authentication from the captive portal with the Active Directory domain.
- g. Select the **Enabled** check box and click **OK** to save your settings and then **Commit** the configuration.
- 3. In the User-ID Agent window, look at the **Connected** column for the new User-ID Agent configuration. A green icon indicates that communication has been established with the firewall that is serving user mapping data.

If a red icon appears, check traffic logs to determine why communications is not working properly. You can also check to see if any user mapping data has been received by running the operational command **show user ip-user-mapping** to view user mapping information on the dataplane, or **show user ip-user-mapping-mp** to view mappings on the management plane.



Note: When redistributing user mapping information to systems with multiple virtual systems, the following configuration is needed:

- Distribution Firewall—Each virtual system must have the collector configured with a pre-shared key. Refer to "Redistribution tab" on page 288. You also need to allow the User-ID Agent services from Network > Network Profiles > Interface Mgmt.
- Receiving Firewall—The collector name and pre-shared secret is needed in order to decide which virtual system User-ID information should be received.

Setting Up the User-ID Agent

The User-ID Agent interfaces with Active Directory or eDirectory to communicate IP addressto-user name mapping information to the firewall.

The User-ID Agent is available for download from Palo Alto Networks. You can install the agent on one or more Windows PCs on your network to obtain user-specific information. When user identification is configured, the firewall's Application Command Center, App-Scope, and logs all include the user name in addition to the user IP address.

Follow the instructions in this section to install and configure the User-ID Agent.



Note: If the multiple virtual system capability is on, you can configure one or more agents per virtual system. This is useful to separate user identification in support of ISPs or other entities that maintain separate user records.

Installing the User-ID Agent

The system on which the User-ID Agent is installed must be running one of the following operating systems/directory server versions:



Note: The 32-bit or 64-bit versions of the below operating systems are supported.

Host Computers

- Microsoft Windows XP/Vista/7
- Microsoft Windows Server 2003/2008

Directory Servers

- Microsoft Exchange
 - 2003 (6.5)
 - 2007 (8.0)
 - 2010 (14.0)
- Microsoft Active Directory
 - 2003
 - 2003r2
 - 2008
 - 2008r2
- Novell eDirectory Server 8.8



Note: Make sure that you choose the correct installation option for your client operating system (32-bit or 64-bit).

Each PC that is included for user identification must be part of the authentication domain. For machines that are not part of the domain, you can use the captive portal capability to screen users and verify user names and passwords.

Refer to these sections for additional information:

- "Configuring the Firewall for User Identification" on page 286—Describes how to set up the firewall to communicate with the User-ID Agents and support captive portals.
- "Captive Portal Policies" on page 206—Describes how to set up captive portal policies.

To install the User-ID Agent, open the installer file and follow the on-screen instructions.

Configuring the User-ID Agent

To open the User-ID Agent:

1. Choose Start > All Programs > Palo Alto Networks > User-ID Agent.

	ID Agent			
Help				
				Commit Exit
User Identification Setup O Discovery Monitoring Cogs	Agent Status Agent is running		Star	t Stop
	Connected Devices			
	Device Address	Status		
	Connected Servers			
	Connected Servers	Туре	Status	
	Connected Servers Server	Туре	Status	
	Connected Servers Server	Туре	Status	
	Connected Servers	Туре	Status	
	Connected Servers	Туре	Status	
	Connected Servers Server	Туре	Status	
	Connected Servers Server	Туре	Status	
	Connected Servers Server	Туре	Status	

Figure 36. User-ID Agent Window

The window contains the following areas and functions:

- Agent Status—Displays the current status of the User-ID Agent.
- **Connected Devices**—Displays the list of devices that the User-ID Agent is currently connected to with associated status.
- **Connected Servers**—Displays the list of servers that the User-ID Agent is currently connected to with associated type and status.

To configure the User-ID Agent:

- 1. Choose Start > All Programs > Palo Alto Networks > User Identification Agent.
- 2. Click **Setup** to open the configuration window.

					Commit Ex
Identification					
Setup					
Discovery	Setup				^
itoring	Service Logon Account Use	ername		barbara@	
long	Security Log Monitor Frequ	ency (sec.)		1	
Logs	Server Session Read Frequ	uency (sec.)		10	
	Novell eDirectory Query In	terval (sec.)		30	
	Enable WMI Probing			Yes	
	Enable NetBIOS Probing			Yes	
	WMI/NetBIOS Probing Inte	erval (min.)		20	
	Enable User Identification	Timeout		No	
	User Identification Timeout	t (min.)		45	
	User-ID Service TCP Port			5007	
	Enable User-ID XML API			No	
	User-ID XML API TCP Port			5006	(70)
	Edit	ase			>
	Edit Access Control List	ase			
	Edit Access Control List Name	Action	From IP Address	To IP Address	
	Edit Access Control List Name	Action	From IP Address	To IP Address	
	Edit Access Control List Name	Action	From IP Address	To IP Address	
	Edit Access Control List Name	Action	From IP Address	To IP Address	
	Edit Access Control List Name	Action	From IP Address	To IP Address	
	Edit Access Control List Name	Action	From IP Address	To IP Address	
	Edit Access Control List Name	Action	From IP Address	To IP Address	
	Edit Access Control List Name	Action	From IP Address	To IP Address	
	Knyell e Directory Search Bi Edit Access Control List Name	Action	From IP Address	To IP Address	
	Edit Access Control List Name	Action	From IP Address	To IP Address	
	Edit Access Control List Name	Action	From IP Address	To IP Address	
		Action	From IP Address	To IP Address	
		Action	From IP Address	To IP Address	
		Action Delete Clo	From IP Address	To IP Address	
	Edit Access Control List Add Edit	Action	From IP Address	To IP Address	

Figure 37. User Identification Configuration Window

- 3. The top of the window lists the current configuration settings. To modify the settings, click **Edit** just below the configuration summary and specify the following settings:
 - **Authentication**—Specify the user name and password to authenticate for Active Directory, WMI, NetBIOS, or eDirectory.
 - Server Monitor—Specify the frequency in seconds (default 1 second) for security log monitor and server session read (default 10 seconds) for Windows server and the query interval for Novell eDirectory query interval. (default 30 seconds).
 - Client Probing—Select the Enable WMI Probing check box if you want to enable
 WMI probing for each workstation and the Enable NetBIOS Probing check box if you want to enable NetBIOS probing for each workstation. Specify an interval between probes (seconds, default 20). An interval of 0 disables this feature.



Note: For WMI polling to work effectively, the Pan Agent service must be configured with a domain administrator account, and each probed client PC must have a remote administration exception configured in the Windows firewall.

Note: For NetBIOS probing to work effectively, each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled.

- **Cache**—Select the check box to enable timeout for the user ID and group cache, and specify the interval (minutes) after which the timeout occurs. Default 45 minutes.
- **Agent Service**—Specify the TCP ports for the user ID service (default 5007) and the user ID XML API (default 5006). Select the check box to enable use of the API.
- eDirectory—Specify the following settings:
 - > **Search Base**—Specify the starting point or root context for agent queries. Example: dc=domain1, dc=example, dc=com.
 - > **Bind Distinguished Name**—Specify the account to bind to the LDAP server. Example: cn=admin, ou=IT, dc=domain1, dc=example, dc=com.
 - > **Bind Password**—Specify the bind account password. The agent saves the encrypted password in the configuration file.
 - Search Filter—Specify the search query for LDAP entries (default is objectClass=Person).
 - > **Search Interval**—Specify the time interval between consecutive queries from the User-ID Agent (range 1-36000 secs, default 30 secs).
 - Server Domain Prefix—Specify a prefix to uniquely identify the user. Use if there are overlapping name spaces. Example: Different users with the same name from two different directories.
 - > **Use SSL**—Select the check box to use SSL for eDirectory binding. If SSL is not selected, a pop-up window warns that clear text will be used for the login account and password.
 - Verify Server Certificate—Select the check box to verify the eDirectory server certificate when using SSL. Select the Enable Group Cache check box to enable the user-group membership cache. When this check box is selected, the user-group membership is cached, and when started, it first reloads the user-group membership from the cache to speed up the restart process.
- 4. Click **Save** to save the configuration.

The User-ID Agent is restarted if the configuration is saved successfully. You can also click the **OK** button to save the configuration and restart the User-ID Agent. If you do not want to restart the User-ID Agent, click **Cancel** to close the dialog box.



Note: You can exclude a list of user accounts from being mapped by creating an ignore-user list. To do this, you create an ignore_user_list.txt file that is placed in the User-ID Agent folder on the domain server where the agent is installed.

As of PAN-OS 4.1 and later, there is no limit to the number of accounts that can be added to these lists. In PAN-OS 4.0 and earlier, the limit was 100 users.

Discovering Domain Controllers

The list of domain controllers available for domain login can be retrieved via DNS. To display the discovery options, click **Discover** on the side menu. You can perform the following tasks in this window:

- Add a configuration setting that allows an administrator to configure the User-ID agent to automatically discover available domain controllers for event log monitoring. Click Add or Edit in the Servers area and specify a server name, IP address, and type (Microsoft Active Directory, Microsoft Exchange, or Novell eDirectory).
- Specify an access control list for networks. Click **Add** or **Edit** in the **Include/exclude lists of configured networks** area, choose the include or exclude option, and specify the network name and address. You can also clone and then modify an existing entry.
- Click **Auto Discover** to automatically retrieve the list of available domain controllers from DNS and add those to the list of monitored servers.

Monitoring User-ID Agent Operation

To view the list of currently discovered IP address-to-user name mapping information, click **Monitoring** on the side menu. You can search for users, or delete users from the list.

To view log entries for the User-ID agent, click **Logs** on the side menu. From this window you can search for log entries or clear the log.

Uninstalling and Upgrading the User-ID Agent

To uninstall the User-ID Agent, open the Control Panel on the PC, select **Add or Remove Programs**, and remove the program **User Identification Agent**.

If you install a new version of the agent and the installer detects an existing installation on your PC, the installer automatically removes the older version before performing the installation.

We recommend that you back up the config.xml file before upgrading the User-ID Agent.

Setting Up the Terminal Services Agent

The Terminal Server Agent (TS Agent) allows the firewall to support multiple users with the same source IP address by identifying the individual firewall users that the terminal server supports.

The TS Agent monitors the remote user sessions and reserves a different TCP/UDP source port range for each user. After a port range is allocated for the user, the TS Agent provides information to map the source port range to the user name.

In addition, the TS Agent requests that the TCP/UDP transport driver in the terminal server allocate the TS-Agent-specified source port instead of the operating system-determined ephemeral port for outbound TCP/UDP traffic. When the firewall receives the TCP/UDP traffic from the terminal server, it checks the source port and obtains the user ID in the ports-to-user map data for the terminal server.

For information on configuring the firewall for terminal services, refer to "Configuring the Firewall for User Identification" on page 286.

Installing or Upgrading the Terminal Server Agent on the Terminal Server

You can install the TS Agent on the following platforms:

- Microsoft Terminal Services 2003
- Microsoft Terminal Services 2008
- Citrix Metaframe Presentation Server 4.0
- Citrix Metaframe Presentation Server 4.5, Citrix XenApp 5, 6

To install the TS Agent on the terminal server:

- 1. Download and open the installation file.
- 2. The installer first checks for platform compatibility. If the platform is not compatible, an error message is displayed.
- 3. The installer checks whether an existing TS Agent exists on the system. If the installer detects that the TS Agent already exists on the system (you are upgrading the TS Agent), it first uninstalls the agent before running the installer.
 - If you are installing a TS Agent that has a newer driver than the existing installation, the installation wizard prompts you to reboot the system after upgrading in order to use the new driver.
 - If you are installing a TS Agent with the same driver version as the existing installation, you can perform the installation as prompted, and do not need to reboot the system afterwards.
- 4. Follow the installer instructions to specify an installation location and complete the installation.



Note: If you specify a destination folder other than the default one, make sure that you use the same destination when you upgrade the TS Agent in the future. If you do not, the existing configuration will be lost and the default configuration will be used.

5. Following installation, reboot the terminal server, if prompted to do so.

Configuring the Terminal Server Agent on the Terminal Server

To configure the TS Agent on the terminal server:

- 1. Launch the TS Agent application from the **Start** menu.
- 2. The configuration panel opens with **Terminal Server Agent** highlighted on the left side of the window.

^p Configure Monitor	Connection List
	Device IP Connection Status 10.1.7.9 : 57662 Connected
	Device Access Control
	Add
	iGave.

Figure 38. Terminal Server Agent Configuration - Main Panel

The connection list box shows all the Palo Alto Networks devices that connect to the TS Agent. The **Device IP** column shows the device IP and port; and the **Connection Status** column indicates whether the status is Connected, Disconnected, or Connecting. Disconnected items are removed from the **Connection List** box when you close and then reopen the TS Agent configuration window.

3. Select the **Enable Device Access Control List** check box if you want to explicitly list the firewalls that the TS Agent will accept. Add each device IP address and click **Add**. Click **Remove** to delete an address from the list. Click **Save** to save the allow list.

4. Click **Configure** to display the configuration settings.

Terminal Server Agent Configure Ionitor	System Source Port Allocation Range: 1025 - 5000 System Reserved Source Ports: Listening Port: 5009 Source Port Allocation Range: 20000 Reserved Source Ports: Port Allocation Start Size Per User: 400
	Port Allocation Maximum Size Per User: 4000

Figure 39. Terminal Server Agent Configuration - Configure Panel

5. Configure settings as described in the following table, and then click **Save**.



Note: If you enter an incorrect parameter and then attempt to save the configuration, a message is displayed to indicate that the configuration will not be saved unless you modify the parameter correctly.

Table 130. Te	rminal Server	Agent Confi	iguration :	Settings
---------------	---------------	-------------	-------------	----------

Field	Description
System Source Port Allocation Range	Displays the port range for system processes that are not associated with individual users. When a server process opens a socket to send a UDP packet or set up a TCP connection, it must obtain a source port from the server operating system. The server automatically allocates a source port (an ephemeral port) for this process. Format is low-high (default 1025- 5000).
	The system port range must not overlap with the Source Port Allocation Range. If they overlap, an application using the system ephemeral source port range could mistakenly be identified as a particular user if the operation system allocated source port falls within the port range allocated for that user.
	Note: Modifying this value requires a Registry change and cannot be done from this panel.
System Reserved Source Ports	Displays the port or ports to be excluded from the operating system source port allocation (because other server processes may use them).
	You can enter a range: <i>low-high</i> (no default).
	Note: Modifying this value requires a Registry change and cannot be done from this panel.

Field	Description
Listening Port	Enter the port on which the terminal server will listen for communications from Palo Alto Networks firewalls (default 5009).
Source Port Allocation	Enter a port allocation range for user sessions.
Range	This setting controls the source port allocation for processes belonging to remote users (default 20000-39999). If a port allocation request comes from system services that cannot be identified as a particular user process, the TS Agent lets the system allocate the source port from the system port range, excluding system reserved source ports.
	Note: Make sure that this port range does not overlap with the System Source Port Allocation Range. If they overlap, an application using the system ephemeral source port range could mistakenly be identified as a particular user if the operation system allocated source port falls within the port range allocated for that user.
Reserved Source Ports	Enter the reserved port allocation range for user sessions. These ports are unavailable for user sessions.
	To include multiple ranges, use commas with no spaces, as in this example: 2000-3000,3500,4000-5000.
	Format is <i>low-high</i> (no default).
Port Allocation Start Size Per User	Enter the number of ports that the TS Agent will first allocate when the remote user logs in (default 200).
	When the remote user logs on, the TS Agent allocates a port range from the Source Port Allocation Range with this specified size. This allows identification of user traffic based on the source port.
Port Allocation Maximum Size Per User	Enter the maximum number of ports that the TS Agent can allocate for a remote user session (default 200).
	If the Port Allocation Start Size Per User setting is not sufficient for the user session, the TS Agent will allocate additional ports up to this maximum.
Fail port binding when	Select the check box as appropriate:
available ports are used up	• If the check box is selected (default), the port request from this user's application fails if the user application has used all available ports. As a result, the application may fail to send traffic.
	• If the check box is not selected, the port request from this user's applica- tion is granted from the System Source Port Allocation Range even if the user application has used all the available ports. The application can send traffic; however, the user ID of the traffic is unknown.

Table 130. Terminal Server Agent Configuration Settings (Continued)

6. Click **Monitor** to display the port allocation information for all terminal server users.

Ally Terminal Server Agent Configure Monitor	Refresh Ports Count	Refresh Intervat	seconds
	User Name	Ports Range	Ports Count
	yyj\test1 yyj\test2	20000-20399 20400-20799	0 4
	User Name: yyyWest	2	ľ
	Ports Ranger 20400-	20799	-
	Ports Count 4		

Figure 40. Terminal Server Agent Configuration - Monitor Panel

7. View the displayed information. For a description of the type of information displayed, refer to the following table.

Field	Description
User Name	Displays the user name.
Ports Range	Displays the current allocated source ports for this user. Multiple ranges are separated by commas (for example, "20400-20799, 20500-20599").
	The size of the port ranges is limited by the "Port Allocation Start Size Per User" and "Port Allocation Maximum Size Per User" configuration parameters, as described in Table 130.
Ports Count	Indicates the number of ports in use.

Table 131. Terminal Server Agent Monitor Information

8. Click the **Refresh Ports Count** button to update the **Ports Count** field manually, or select the **Refresh Interval** check box and configure a refresh interval to update this field automatically.

The following table lists the menu options available in the TS Agent application window.

Field	Description
Configure	Open the Configuration panel.
Monitor	Open the Monitor panel.
Restart Service	Restart the TS Agent service. This option is not normally required and is reserved for troubleshooting.
Show Logs	Display the troubleshooting log.

Table 132. Terminal Server Agent Menu Options

Field	Description
Debug	Select debugging options (None, Error, Information, Debug, or Verbose).
Exit	Quit the TS Agent application.
Help	Display TS Agent version information.

Table 132. Terminal Server Agent Menu Options (Continued)

Uninstalling the Terminal Server Agent on the Terminal Server

To uninstall the TS Agent, use the **Add/Remove Programs** control panel on the server. Remove the "Terminal Server Agent" application. You must reboot the system to complete the uninstallation. Setting Up the Terminal Services Agent

Chapter 8 Configuring IPSec Tunnels

This chapter describes basic virtual private network (VPN) technology and provides details on configuring IP Security (IPSec) VPNs on Palo Alto Networks firewalls. Refer to the following sections:

- "Virtual Private Networks" in the next section
- "IPSec and IKE" on page 311
- "Setting Up IPSec VPNs" on page 313
- "Sample VPN Configuration" on page 320
- "GlobalProtect Large Scale VPN Deployment" on page 323

Virtual Private Networks

Virtual private networks (VPNs) allow systems to connect securely over a public network as if they were connecting over a local area network (LAN). The IP Security (IPSec) set of protocols is used to set up a secure tunnel for the VPN traffic, and the private information in the TCP/IP packets is encrypted when sent through the IPSec tunnel. The Palo Alto Networks firewall supports IPv4 tunnels as well as the option to enable IPv6 in the tunnel configuration. This option allows you to route IPv6 traffic over the IPv4 tunnel to provide confidentiality between IPv6 networks when IPv6 WAN connectivity is not available.

The Palo Alto Networks GlobalProtect Large Scale VPN feature provides a greatly simplified mechanism to roll out a hub and spoke VPN that utilizes certificates for authentication and automatically refreshes the authentication credentials periodically. Refer to "GlobalProtect Large Scale VPN Deployment" on page 323.



Note: In addition to IPSec VPNs, the firewall also supports Secure Sockets Layer (SSL) VPNs, which allow remote users to establish VPN connections through the firewall. Refer to **Chapter 9**, "**Configuring GlobalProtect**" for more information.

The following figure shows a standard IPSec tunnel between two devices. The configuration can include a tunnel monitor on each side to alert the device administrator of tunnel failure and provide automatic failover. Tunnel monitors are useful if you want to be able to provide failover of IPSec traffic to another interface.





You can configure route-based VPNs to connect Palo Alto Networks firewalls at central and remote sites or to connect Palo Alto Networks firewalls with third-party security devices at other locations. With route-based VPNs, the firewall makes a routing decision based on the destination IP address. If traffic is routed to a specific destination through a VPN tunnel, then it is encrypted as VPN traffic. It is not necessary to define special rules or to make explicit reference to a VPN tunnel; routing and encryption decisions are determined only by the destination IP address.

The firewall can also interoperate with third-party policy-based VPN devices. To connect with a policy-based VPN, configure the Proxy ID for the tunnel. If multiple phase 2 tunnels are required, configure different Proxy IDs on each. Refer to "Setting Up IPSec Tunnels" on page 315.

For the IPSec connection between the firewalls, the full IP packet (header and payload) is embedded in another IP payload, and a new header is applied. The new header uses the IP address of the outgoing firewall interface as the source IP address and the incoming firewall interface at the far end of the tunnel as the destination IP address. When the packet reaches the firewall at the far end of the tunnel, the original packet is decrypted and sent to the actual destination host.

IPSec Security Associations (SAs) are defined at each end of the IPSec tunnel to apply all of the parameters that are required for secure transmission, including the security parameter index (SPI), security protocol, cryptographic keys, and the destination IP address. Encryption, data authentication, data integrity, and endpoint authentication are provided by IPSec SAs.

VPN Tunnels

To set up VPNs, it is important to understand your network topology and be able to determine the required number of tunnels. For example:

- A single VPN tunnel may be sufficient for connection between a single central site and a remote site.
- Connections between a central site and multiple remote sites require VPN tunnels for each central remote site pair.

Each tunnel is bound to a tunnel interface. It is necessary to assign the tunnel interface to the same virtual router as the incoming (clear text) traffic. In this way, when a packet comes to the firewall, the route lookup function can determine the appropriate tunnel to use. The tunnel interface appears to the system as a normal interface, and the existing routing infrastructure can be applied.

Each tunnel interface can have a maximum of 10 IPSec tunnels. This allows you to set up IPSec tunnels for individual networks that are all associated with the same tunnel interface on the firewall.

IPSec and IKE

There are two ways to secure IPSec VPN tunnels:

- Configure the tunnel using manual security keys. This method is not recommended.
- Generate keys using Internet Key Exchange (IKE)

The same method must be applied to both ends of the IPSec tunnel. In the case of manual keys, the same key is entered at both ends; in the case of IKE, the same methods and attributes are applied at both ends.

IKE provides a standard mechanism for generating and maintaining security keys:

- **Identification**—The identification process involves recognition of the peers at both ends of the IPSec tunnel. Each peer is identified by IP address or peer ID (contained in the payload of the IP packet). The firewall or other security device at each end of the tunnel adds the identification of the peer at the other end into its local configuration.
- Authentication—There are two types of authentication methods: pre-shared key and PKI. Currently only the pre-shared key method is supported by Palo Alto Networks firewalls.

The firewall supports definition of IKE gateways, which specify the configuration information necessary to perform IKE protocol negotiation with peer gateways.

IKE configuration options include Diffie-Hellman Group for key agreement, Encryption algorithm, and hash for message authentication.

IPSec and IKE Crypto Profiles

Crypto profiles are related to standard proposal fields in IKE negotiation.

- IKE Phase-1 authenticates the firewalls to each other and sets up a secure control channel. It uses the IKE-crypto profile for IKE SA negotiation.
- IKE Phase-2 is the negotiation, through the Phase 1 SA, of an actual tunnel for traffic between networks behind the respective firewalls. It uses the IPSec crypto profile for IPSec SA negotiation.

You can define IPSec and IKE crypto profiles that determine the protocols and algorithms used to negotiate the IPSec and IKE SAs.

Options for IKE SA:

- **Diffie-Hellman (DH) Group**—Select DH groups to use when generating public keys for IKE.
- Encryption—Select encryption algorithms.
- Hash Algorithm—Select hash algorithms.
- Lifetime—Specify the length of time that the negotiated key will stay effective.

Options for IPSec SA:

- Encapsulating Security Payload (ESP)—Select options for authentication, data integrity, confidentiality, and encryption.
- Authentication Header (AH)—Select options for authentication and data integrity. This option is not generally used.
- **Perfect Forward Security (PFS) Diffie-Hellman (DH) group**—Select DH groups to use in generating independent keys for IPSec.
- Lifetime—Specify the length of time that the negotiated key will stay effective.

For details on the specific protocols and algorithms supported for IPSec and IKE crypto profiles, refer to "Defining IKE Crypto Profiles" on page 318 and "Defining IPSec Crypto Profiles" on page 319.

Setting Up IPSec VPNs

This section describes the multi-step process involved in setting up IPSec VPN tunnels. For detailed instructions, refer to the specified sections in this guide. For a sample configuration, refer to "Sample VPN Configuration" on page 320.



Note: Before you begin, make sure that your Ethernet interfaces, virtual routers, and zones are configured properly. Refer to "Firewall Interfaces" on page 127, "Virtual Routers and Routing Protocols" on page 153, and "Defining Security Zones" on page 151.

To set up IPSec VPNs:

- 1. Plan the network topology and determine the required number of tunnels.
- 2. Define IKE gateways with the configuration information for IKE protocol negotiation with peer gateways. Refer to "Defining IKE Gateways" on page 314.
- 3. Configure the protocols and algorithms for identification, authentication, and encryption in VPN tunnels using IKE SA negotiation:
 - For IKEv1 Phase-1, refer to "Setting Up IPSec Tunnels" on page 315.
 - For IKEv1 Phase-2, refer to "Defining IPSec Crypto Profiles" on page 319.
- 4. Configure the parameters that are needed to establish IPSec VPN tunnels. Refer to "Setting Up IPSec Tunnels" on page 315.
- 5. Specify how the firewall will monitor the IPSec tunnels. Refer to "Viewing IPSec Tunnel Status on the Firewall" on page 319.
- 6. Set up static routes or assign routing protocols to redirect traffic into the newly established tunnels. The Border Gateway Protocol, Routing Information Protocol (RIP), and Open Shortest Path First (OSPF) options are supported; you can enable these protocols on the tunnel interface. Refer to "Virtual Routers and Routing Protocols" on page 153.
- 7. Set security policies to filter and inspect the traffic as described in "Security Policies" on page 187. Define the source and destination zones and specify the policy attributes as follows:
 - **Outgoing traffic entering the tunnel**—For source, use the clear text zone. For destination, use the tunnel interface zone.
 - **Incoming traffic egressing the tunnel**—For source, use the tunnel interface zone. For destination, use the clear text zone.

After defining the zones for the security policy rule, set the source and destination addresses, applications, services, and security profiles to control access across the VPN tunnel.



Note: If the tunnel interface is in the same zone as the destination clear text traffic, the same policies will be used for VPN and clear text traffic. Ideally, you will want to put the tunnel interface in a separate zone, so tunneled traffic can use different policies.

When these tasks are complete, the tunnel is ready for use. Traffic destined for the addresses defined for the tunnels is automatically routed properly and encrypted as VPN traffic based on the specific destination route added to the routing table.



Note: Without matching security rules, VPN traffic will be dropped by the firewall, when a security rule is required.

The IKE protocol will be triggered when necessary (for example, when traffic is routed to an IPSec tunnel with no keys or expired keys).

If there is a deny rule at the end of the security rulebase, intra-zone traffic is blocked unless otherwise allowed. Rules to allow IKE and IPSec applications must be explicitly included above the deny rule.

Defining IKE Gateways

Network > Network Profiles > IKE Gateways

Use the **IKE Gateways** page to define gateways that include the configuration information necessary to perform IKE protocol negotiation with peer gateways.

Field	Description
Name	Enter a name to identify the gateway (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Interface	Specify the outgoing firewall interface.
Local IP Address	Select the IP address for the local interface that is the endpoint of the tunnel.
Peer Type	Static IP address or dynamic option for the peer on the far end of the tunnel.
Peer IP Address	If the Static option is selected for peer type, specify the IP address for the peer on the far end of the tunnel.
Pre-Shared Key Confirm Pre-Shared Key	Enter a security key to use for authentication across the tunnel. Applies for static and dynamic peer types.

Table 133. IKE Gateway Settings

Note: The following advanced fields are visible if you select the check box for the **Show advanced Phase 1** *options* link.

Local Identification	Choose from the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), KEYID (binary format ID string in HEX). If no value is specified, the local IP address will be used as the local identification value.
Peer Identification	Choose from the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), KEYID (binary format ID string in HEX). If no value is specified, the peer IP address will be used as the peer identification value.
Exchange Mode	Choose auto, aggressive, or main.

Field	Description
IKE Crypto Profile	Select an existing profile or keep the default profile.
Enable Passive Mode	Select to have the firewall respond only to IKE connections and never initiate them.
Enable NAT Traversal	Select to have UDP encapsulation used on IKE and UDP protocols, enabling them to pass through intermediate NAT devices.
	NAT traversal is used when NAT addressing is in place between the IPSec VPN terminating points.
Dead Peer Detection	Select the check box to enable and enter an interval (2 - 100 seconds) and delay before retrying (2 - 100 seconds). Dead peer detection identifies inactive or unavailable IKE peers through ICMP ping and can help restore resources that are lost when a peer is unavailable.

Table 133. IKE Gateway Settings (Continued)



Note: When a device is set to use the **auto** exchange mode, it can accept both main mode and aggressive mode negotiation requests; however, whenever possible, it initiates negotiation and allows exchanges in main mode.

You must configure the peer device with the matching exchange mode to allow it to accept negotiation requests initiated from the first device.

Setting Up IPSec Tunnels

Network > IPSec Tunnels

Use the **IPSec Tunnels** page to set up the parameters to establish IPSec VPN tunnels between firewalls.

Field	Description
General Tab	
Name	Enter a name to identify the tunnel (up to 63 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
	The 63 character limit for this field includes the tunnel name in addition to the Proxy ID, which is separated by a colon character.
Tunnel Interface	Select an existing tunnel interface, or click New Tunnel Interface to create a new tunnel interface. For information on creating a tunnel interface, refer to "Configuring Tunnel Interfaces" on page 147.
Туре	Select whether to use an automatically generated or manually entered security key. Auto key is recommended.

Table 134. IPSec Tunnel Settings

Field	Description
Auto Key	If you choose Auto Key , specify the following: • IKE Gateway —Refer to "Defining IKE Gateways" on page 314 for
	 IPSec Crypto Profile—Select an existing profile or keep the default profile. To define a new profile, click New and follow the instructions in "Defining IPSec Crypto Profiles" on page 319.
	Advanced
	• Enable Replay Protection—Select this option to protect against replay attacks.
	• Copy TOS Header —Copy the (Type of Service) TOS header from the inner IP header to the outer IP header of the encapsulated packets in order to preserve the original TOS information.
	• Tunnel Monitor —Select this option to alert the device administrator of tunnel failures and to provide automatic failover to another interface. Note that you need to assign an IP address to the tunnel interface for monitoring.
	 Destination IP—Specify an IP address on the other side of the tunnel that the tunnel monitor will use to determine if the tunnel is working properly.
	 Profile—Select an existing profile that will determine the actions that are taken if the tunnel fails. If the action specified in the monitor profile is wait-recover, the firewall will continue to use the tunnel interface in routing decisions as if the tunnel remained active. If the fail-over action is used, the firewall will disable the tunnel interface, thereby disabling any routes in the routing table that use the interface. For more information, see "Defining Monitor Profiles" on page 177.
Manual Key	If you choose Manual Key, specify the following:
,	• Local SPI—Specify the local security parameter index (SPI) for packet traversal from the local firewall to the peer. SPI is a hexadecimal index that is added to the header for IPSec tunneling to assist in differenti-ating between IPSec traffic flows.
	• Interface —Select the interface that is the tunnel endpoint.
	• Local Address—Select the IP address for the local interface that is the endpoint of the tunnel.
	• Remote SPI —Specify the remote security parameter index (SPI) for packet traversal from the remote firewall to the peer.
	• Protocol —Choose the protocol for traffic through the tunnel (ESP or AH).
	 Authentication—Choose the authentication type for tunnel access (SHA1, SHA256, SHA384, SHA512, MD5, or None).
	• Key/Confirm Key—Enter and confirm an authentication key.
	 Encryption—Choose an encryption option for tunnel traffic (3des, aes128, aes192, aes256, aes128ccm16, or null [no encryption]).
	• Key/Confirm Key—Enter and confirm an encryption key.

 Table 134.
 IPSec Tunnel Settings (Continued)

Field	Description
GlobalProtect Satellite	If you choose GlobalProtect Satellite , specify the following:
	For an overview of the concept of GlobalProtect Satellites, refer to "GlobalProtect Large Scale VPN Deployment" on page 323.
	• Name—Enter a name to identify the tunnel (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
	• Tunnel Interface —Select an existing tunnel interface, or click New Tunnel Interface. For information on creating a tunnel interface, refer to "Configuring Tunnel Interfaces" on page 147.
	• Portal Address—Enter the IP address of the GlobalProtect Portal.
	• Interface —Select the interface from the drop-down that is the egress interface to reach the GlobalProtect Portal.
	• Local IP Address—Enter the IP address of the egress interface that connects to the GlobalProtect Portal.
	Advanced Options
	• Publish all static and connected routes to Gateway —Select this option to publish all routes from the satellite device to the GlobalProtect Gateway in which this satellite is connected.
	• Subnet —Click Add to manually add local subnets for the satellite loca- tion. If other satellites are using the same subnet information, you must NAT all traffic to the tunnel interface IP. Also, the satellite must not share routes in this case, so all routing will be done through the tunnel IP.
	• External Certificate Authority—Select this option if you will use an external CA to manage certificates. Once you have your certificates generated, you will need to import them into the device and select the Local Certificate and the Certificate Profile to be used.
Proxy ID Tab	
Proxy ID	Click Add and enter a name to identify the proxy.
Local	Enter an IP address or subnet in the format <i>ip_address/mask</i> (for example, 10.1.2.1/24).
Remote	If required by the peer, enter an IP address or subnet in the format $ip_address/mask$ (for example, 10.1.1.1/24).
Protocol	Specify the protocol and port numbers for the local and remote ports:
	 Number—Specify the protocol number (used for interoperability with third-party devices).
	• Any—Allow TCP and/or UDP traffic.
	• TCP—Specify the local and remote TCP port numbers.
	• UDP —Specify the local and remote UDP port numbers.
	<i>Note:</i> Each configured proxy ID will count towards the IPSec VPN tunnel capacity of the firewall.

 Table 134.
 IPSec Tunnel Settings (Continued)

Important items to consider when configuring IPSec VPNs

Keep the following in mind when configuring IPSec VPNs:

• There must be a route to the remote network that is being tunneled.

- Pre-shared keys may be entered incorrectly on one of the devices. Pre-shared keys must always match.
- Phase 1 negotiation mode (aggressive/main) may not match on the devices. The negotiation mode must always match.
- A common misconfiguration is to enable perfect forward secrecy on only one side. It must be enabled on both sides.
- If the dynamic routing protocols advertise routes to public IP addresses through the IPSec tunnel, the device establishing the tunnel may attempt phase 1 negotiation with the destination set to the public IP rather than the endpoint of the IPSec tunnel. As a result, the connection is never created and routing fails. To address this problem, ensure that only private IP addresses route through the tunnel and that no public IP addresses or default routes exist in the routing table that points to the tunnel.
- A Proxy ID may be improperly entered for the device at the far end of the IPSec tunnel. This can occur because some vendors generate a default Proxy ID for IPSec communications that is not easily identified by the end user.

Defining IKE Crypto Profiles

Network > Network Profiles > IKE Crypto

Use the **IKE Crypto Profiles** page to specify protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPSec SA negotiation (IKEv1 Phase-1). Refer to "Virtual Private Networks" on page 310 for more information.

To change the ordering in which an algorithm or group is listed, select the item and then click the **Move Up** or **Move Down** icon. The ordering determines the first choice when settings are negotiated with a remote peer. The setting at the top of the list is attempted first, continuing down the list until an attempt is successful.

Field	Description
DH Group	Specify the priority for Diffie-Hellman (DH) groups. Click Add and select groups. For highest security, select an item and then click the Move Up or Move Down icon to move the groups with higher numeric identifiers to the top of the list. For example, move group14 above group2 .
Authentication	Specify the priority for hash algorithms. Click Add and select algorithms (md5, sha1, sha256, sha384, or sha512). For highest security, use the arrows to move sha1 to the top of the list.
Encryption	Select the check boxes for the desired Encapsulating Security Payload (ESP) authentication options. Click Add and select algorithms (aes256, aes192, aes128, or 3des). For highest security, select an item and then click the Move Up or Move Down icon to change the order to the following: aes256 , aes192 , aes128 , 3des .
Lifetime	Select units and enter the length of time that the negotiated key will stay effective.

Defining IPSec Crypto Profiles

▶ Network > Network Profiles > IPSec Crypto

Use the **IPSec Crypto Profiles** page to specify protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPSec SA negotiation (IKEv1 Phase-2). Refer to "Virtual Private Networks" on page 310 for more information.

Field	Description
Name	Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
IPSec Protocol	Choose an option from the drop-down list.
	ESP:
	• Click Add under Encryption and select the desired ESP encryption algorithms. For highest security, use the arrows to change the order to the following: 3des, aes128, aes192, aes256, or aes128ccm16.
	• Click Add under Authentication and select the desired ESP authentica- tion algorithms (md5, sha1, sha256, sha384, sha512, or none).
	AH:
	• Click Add under Authentication and select the desired AH authentica- tion algorithms (md5, sha1, sha256, sha384, or sha512).
DH Group	Select the DH group. For highest security, choose the group with the highest identifier.
Lifetime	Select units and enter the length of time that the negotiated key will stay effective. The default is 1 hour.
Lifesize	Select optional units and enter the amount of data that the key can use for encryption.

Table 136. IPSec Crypto Profile Settings

To change the ordering in which an algorithm or group is listed, select an item and then click the **Move Up** or **Move Down** icon to change the order. The listed order determines the order in which the algorithms are applied and can affect tunnel performance.

Viewing IPSec Tunnel Status on the Firewall

► Network > IPSec Tunnels

To view the status of currently defined IPSec VPN tunnels, open the **IPSec Tunnels** page. The following status information is reported on the page:

- **Tunnel Status (first status column)**—Green indicates an IPSec SA tunnel. Red indicates that IPSec SA is not available or has expired.
- **IKE Gateway Status**—Green indicates a valid IKE phase-1 SA. Red indicates that IKE phase-1 SA is not available or has expired.
- **Tunnel Interface Status**—Green indicates that the tunnel interface is up (because tunnel monitor is disabled, or because tunnel monitor status is UP). Red indicates that the tunnel interface is down, because the tunnel monitor is enabled and the status is down.

Sample VPN Configuration

This section describes a sample VPN configuration. In this sample, a branch office is connected with a headquarters office and branch office users are allowed to access a central server farm.

Refer to the following topics:

- "Existing Topology" in the next section
- "New Topology" on page 321
- "Configure the VPN Connection" on page 321
- "VPN Connectivity Troubleshooting" on page 322

Existing Topology

Headquarters:

- Firewall public IP 61.1.1.1, on interface ethernet1/1, which is in zone "ISP", virtual-router "HQ"
- Server farm network is 10.100.0.0/16, connected through interface ethernet1/5 (IP 10.100.0.1), which is on zone "server", virtual-router "HQ"

Branch office:

- Firewall public IP is 202.101.1.1, on interface ethernet1/2, which is in zone "ISP-branch", virtual-router "branch"
- A PC network of 192.168.20.0/24, connected through interface ethernet1/10, which is on zone "branch-office", virtual-router "branch" (same as ethernet1/2)
- Security policy to allow traffic from zone "branch-office" to zone "ISP-branch" for Internet access from the PC network

The following figure shows the existing topology.



Figure 42. Sample VPN Configuration - Existing Topology

New Topology

Headquarters:

- Create a new security zone "branch-vpn."
- Add a tunnel interface tunnel.1 to zone "branch-vpn" and assign an IP address from a private range (for example, 172.254.254.1/24)
- Add a static route to direct traffic to 192.168.20.0/24 (the branch office network) to the tunnel interface tunnel.1.
- Add a security policy to allow traffic from zone "branch-vpn" to zone "server."

Branch office:

- Create a new security zone "central-vpn."
- Add a tunnel interface tunnel.2 to zone "central-vpn" and assign an IP address from private range (for example, 172.254.254.20/24).
- Add a static route to direct traffic to 10.100.0.0/16 (the server farm network) to the tunnel interface tunnel.2.
- Add a security policy to allow traffic from zone "branch-office" to zone "central-vpn."

The following figure shows the tunnel information for the new topology.





Configure the VPN Connection

Headquarters:

- Create an IKE gateway "branch-1-gw" with these parameters:
 - Peer-address: dynamic (or 202.101.1.1)
 - Local-address: ethernet1/1
 - Peer-ID: type is FQDN: branch1.my.domain
 - Authentication: pre-shared-key newvpn
 - Protocol: keep default values

- Create an IPSec tunnel "branch-1-vpn" with these parameters:
 - ike-gateway-profile: branch-1-gw
 - ipsec-crypto-profile: leave as default
 - Tunnel interface: **bind with tunnel.1**
- On servers in the server farm, check the routing table and verify that the destination 192.168.20.0/24 is reachable through 10.100.0.1.

Branch office:

- Create an IKE gateway "central-gw" with these parameters:
 - Peer-address: 61.1.1.1
 - Local-address: ethernet1/2
 - Local-ID: type is FQDN: **branch1.my.domain**
 - Authentication: pre-shared-key newvpn
 - Protocol: keep default values
- Create an IPSec tunnel "central-vpn" with these parameters:
 - ike-gateway-profile: central-gw
 - ipsec-crypto-profile: leave as default
 - Tunnel interface: bind with **tunnel.2**

Configuration Notes:

- If 202.101.1.1 is set as the peer-address parameter in "branch-1-gw" on the central site, setting the local-id and peer-id parameters becomes unnecessary (the field can be left empty). Note that treatment of these two parameters must be the same, because these two fields are matched during IKE negotiation.
- The proxy-id is left empty for route-based VPNs such as this.

After configuring the parameters and committing the configuration, the new VPN should work. If connectivity issues arise, refer to "VPN Connectivity Troubleshooting" in the next section.

VPN Connectivity Troubleshooting



Note: The parameter values in this section refer to the sample configuration. Refer to "Configure the VPN Connection" on page 321.

To troubleshoot issues regarding VPN connectivity:

- 1. Double check configurations on both sites.
- 2. Use the **ping** utility to verify connectivity between the central and branch offices (202.101.1.1 and 61.1.1.1). For this to work, there must be a management profile on the interface that allows **ping**.

- 3. Use the **ping** utility to verify connectivity between the server farm and the central firewall (ethernet1/5). For this to work, there must be a management profile on the interface that allows **ping**.
- 4. Use the **ping** utility to verify connectivity between the branch network and the branch firewall interface (ethernet1/10). For this to work, there must be a management profile on the interface that allows **ping**.
- 5. On the branch-office site, use the CLI commands **test vpn ike-sa gateway central-gw** and **show vpn ike-sa gateway central-gw** to verify that IKE phase-1 SA can be created from the branch office.
- 6. On the central site, use the CLI command **show vpn ike-sa gateway branch-1-gw** to verify that IKE phase-1 SA can be created from the branch office.
- 7. On the branch office site, use the CLI command **test vpn ipsec-sa tunnel central-vpn** and **show vpn ipsec-sa tunnel central-vpn** to verify that IKE phase-2 SA can be created from the branch office.
- 8. On the central site, use the CLI command **show vpn ipsec-sa tunnel branch-1-vpn** to verify that IKE phase-2 SA can be created from the branch office.
- 9. Check the server routing table in the server farm. The destination 192.168.20.0/24 must be reachable through the central firewall's ethernet1/5 interface IP address.
- 10. To check the route setting, run the **traceroute** command from any PC in the branch office network, where the destination is one of servers in the server farm.
- 11. Run the **ping** utility from any PC in the branch office network, where the destination is one of servers in the server farm. Check the encryption and decryption counters shown in the output of the **show vpn flow** CLI command. Verify that these counters are incrementing and that none of the error counters are incrementing.
- 12. Examine the detailed error messages for IKE negotiation in the syslog or use the **debug ike pcap** command to capture IKE packets in PCAP format.

GlobalProtect Large Scale VPN Deployment

This section contains the following topics:

- "Overview" in the next section
- "Deploying a Large Scale VPN Network" on page 324
- "Dynamic Routing Protocols and Large Scale VPNs" on page 333
- "Backing up a GlobalProtect Portal" on page 334

Overview

Deploying large scale VPN networks can be a very time consuming and complicated process. Challenges include planning, authentication requirements, tunnel configuration, route planning, and the provisioning and the decommissioning of components and devices. With Palo Alto Networks firewalls, this process has been greatly simplified by taking advantage of the deployment and management methods used in the GlobalProtect VPN feature, which has previously only been used for remote access for client PCs.

The illustration in Figure 44 shows that each satellite connects to the portal to obtain a certificate for authentication and then downloads an initial VPN configuration. Once the initial configuration is completed, the satellite device establishes a VPN to all configured gateways defined on the portal using the same trusted certificate, creating a hub-and-spoke VPN network. Network and routing information is then shared between the gateway and satellite devices to create multiple paths to ensure that connectivity between the corporate and branch offices are always maintained.



Figure 44. Large Scale VPN Deployment

Deploying a Large Scale VPN Network

When deploying a large scale VPN network using Palo Alto Networks firewalls and GlobalProtect, you first configure the GlobalProtect Gateway(s), GlobalProtect Portal, and certificates that will be used for authentication of satellite devices. Once the portal and gateway devices are configured, you simply establish WAN connectivity from the satellite device, provision a tunnel interface, enter the host name of the Internet accessible GlobalProtect Portal interface and the satellite device configured gateway devices and utilize tunnel monitoring to discover failures and initiate a failover. Once the portal and gateways are configured, you can easily add additional satellite devices.

It is important to keep a couple things in mind when deploying multiple branch offices:
- If you have branch offices that will use duplicate subnets, the satellite device must NAT all traffic to a unique IP or subnet. One of the simplest options is to source NAT all satellite-initiated traffic to the tunnel interface IP. When specifying the tunnel interface as a source NAT translated address, the firewall will use the IP address assigned to the tunnel interface by the highest priority gateway. If multiple gateways service different networks, it is recommended that you provision a unique IP address or subnet for each satellite that will be routable from all gateway devices. When duplicate subnets are in use, be sure that the satellites are not permitted to share connected routes with the gateways, otherwise routing issues will occur.
- Satellites must use tunnel monitoring (utilizing a configuration provided by the gateways) in order to recognize when a gateway fails. This also ensures that tunnels are maintained when there is no traffic flowing to the gateway.

The following section describes the components and basic steps needed to deploy a large scale VPN network using GlobalProtect.

- "Certificates and the OCSP Responder" on page 325
- "Global Protect Gateway Configuration" on page 328
- "GlobalProtect Portal Configuration" on page 330
- "GlobalProtect Satellite Configuration" on page 332
- "Dynamic Routing Protocols and Large Scale VPNs" on page 333

Certificates and the OCSP Responder

Certificates are used in the large scale VPN environment to provide a mutual authentication scheme for the satellite devices and to automate deployment to reduce the steps needed to prepare the satellite devices.

The GlobalProtect Portal is configured as a root CA that will sign certificates for authentication for the satellite and gateway devices. The gateway uses a certificate profile with the portal's root CA in order to allow satellites to authenticate to the gateways. To add a new satellite device to the VPN network, you simply add the device serial number to the portal configuration or create a username and password that the satellite administrator will use to start the initial configuration. On the satellite, you enter the portal address in the IPSec Tunnels page and the satellite will connect to the portal and will send a certificate signing request (CSR). The portal then signs and returns the client certificate to the satellite as well as the root CA certificate to allow communication.

In the case of the gateway certificates, the portal will create a server certificate that is imported to the gateway, and the portal's root CA certificate is defined in the certificate profile on the gateway. This will allow all satellites to authenticate to the defined gateways.

To manage certificates in the environment, you create an Online Certificate Status Protocol (OCSP) profile on the portal that will be used to manage revocation status for local certificates issued to the satellite and gateway devices. On each satellite and gateway device, you will configure the OCSP option to point to the portal.



Note: You can also use an external root certificate authority (CA) for the GlobalProtect VPN deployment.

For information on certificates, refer to "Importing, Exporting and Generating Security Certificates" on page 86. For information on the OCSP responder, refer to "OCSP Responder" on page 90.

Configure the OCSP Responder:

The OCSP responder needs to be configured first since it will be used when creating certificates for the satellite and gateway devices.

- 1. Navigate to **Device > Certificate Management > OCSP Responder**.
- 2. Click **Add** and enter a name for the OCSP responder.
- 3. Enter the **Host Name** of the portal.

Configure an OCSP Responder Management Profile:

This profile is used to allow the satellite and gateway devices to connect to the portal for OCSP requests over HTTP.

- 1. Navigate to Network > Network Profiles > Interface Mgmt.
- 2. Click **Add** and name the profile.
- 3. Select the HTTP OCSP check box. You may also want to enable ping for testing purposes.
- 4. Click **OK** to save your changes.
- 5. Assign the new management profile to the ingress interface of the portal by navigating to **Network > Interfaces**.
- 6. Select the interface that will be used as the ingress for the VPN devices and click the **Advanced** tab.
- 7. From the **Other Info** tab and the **Management Profile** drop-down, select the OCSP management profile that you created.
- 8. Click **Commit** to activate your changes.

Generate a root Certificate Authority (CA) certificate:

The portal root CA will sign and issue certificates for the satellite and gateway devices to provide mutual authentication for devices that are part of the large scale VPN network. The gateway will also have a certificate profile with the portal's root CA to allow authentication from the satellite devices.

- 1. Navigate to **Device > Certificate Management > Certificates**.
- 2. Click **Generate** and name the certificate.

- 3. Enter the **Common Name**, which is the IP or FQDN that will appear on the certificate. The common name is the FQDN or IP address of the interface where satellites will connect.
- 4. Check the **Certificate Authority** box.
- 5. In the **OCSP Responder** field, enter the OCSP responder name that you created earlier.

Generate a server certificate for the portal:

This certificate is used to authenticate satellite devices and to allow OCSP queries from the satellites and gateways.

- 1. On the GlobalProtect Portal, navigate to **Device > Certificate Management > Certificates**.
- 2. Click **Generate** and name the certificate and enter a common name. **Important:** The common name is the FQDN or IP address of the interface used for the satellite connections and will be unique to the portal.
- 3. In the **Signed By** drop-down, select the root CA of the portal. Set the desired cryptographic settings and other attributes and then click **Generate.**

Generate and import the gateway certificates:

Each gateway that will participate in the VPN network needs to have a server certificate that is signed by the GlobalProtect Portal root CA and the portal root CA certificate must also be imported into each gateway. This will allow the gateways to communicate with the portal and allows satellites devices to connect to the gateway to establish VPN connectivity. In the below steps, you can alternatively export the root CA from the portal to the gateway, and then generate the gateway certificate directly on the gateway using the imported root CA to sign the certificate.

On the GlobalProtect Portal, navigate to **Device > Certificate Management > Certificates**.

- 1. Click **Generate** and name the certificate and enter a common name. **Important:** The common name is the FQDN or IP address of the interface used for the satellite to connect to the gateway and will be unique for each gateway.
- 2. In the **Signed By** drop-down, select the root CA of the portal.
- 3. In the **OCSP Responder**, it is important that you enter the URL of the portal using the following format: **http://IP** address or FQDN/CA/ocsp, for example http:// paloaltonetworks.com/CA/ocsp. Certificates issued by the portal will have this path automatically generated in the certificate.
- 4. Set the desired cryptographic settings and other attributes and then click **Generate**.
- 5. Select the certificate you just created and then click **Export**, select the File Format **Encrypted Private Key and Certificate (PKCS12)**. Enter and confirm a passphrase that will be used when you import the certificate.
- 6. Follow the same steps to export the root CA certificate.
- On the GlobalProtect Gateway, navigate to Device > Certificate Management > Certificates.
- 8. Click **Import**, enter a **Certificate Name** and then click **Browse** and select the gateway server certificate you exported earlier.

- 9. Select the file format **Encrypted Private Key and Certificate (PKCS12)**. Enter and confirm the passphrase you used when you exported the certificate and click **OK**.
- 10. Do the same for the portal root CA.

Configure a certificate profile on the gateway:

The certificate profile is used to authenticate the satellite device to the gateway to establish a the VPN.

- 1. Navigate to **Device > Certificate Management > Certificate Profile**.
- 2. Click **Add** and name the profile. The **Username Field** and **Domain** are optional when only Palo Alto Networks devices will connect to the gateway. You will need to configure these options if you have third-party VPN clients connecting to the gateway, such as with iOS and Android devices.
- 3. In the **CA Certificates** window, click **Add** and in the **CA Certificate** drop-down, select the root CA certificate that you imported from the portal.
- 4. In the **Default OCSP URL** field, enter the URL of the OCSP responder using the IP or FQDN of the portal. It is important that you enter the URL of the portal using the following format: http://IP address or FQDN/CA/ocsp, for example http:// paloaltonetworks.com/CA/ocsp.

If you have another device that can also handle OCSP requests, you can add a certificate for that profile in the **OCSP Verify CA Certificate** drop-down.

- 5. Click **OK**.
- 6. In the **Certificate Profile** window, click the **Use OCSP** check box.
- 7. Click **OK** to save your changes.
- 8. Click **Commit** to active the certificate configuration.

Global Protect Gateway Configuration

The branch office satellite device will connect to the GlobalProtect Gateway to establish VPN connectivity to corporate resources. The satellite will use each gateway that is defined in the GlobalProtect Portal satellite configuration in order to create a hub and spoke VPN network. The gateway devices are also able to maintain multiple tunnels on a single tunnel interface, so you don't need a separate tunnel interface for each satellite device that is deployed. Authentication from the satellite to the gateway is handled by the certificates signed by the portal CA. Once configured and the VPN is established, routing information is also dynamically shared between the satellite and gateway devices.

Configuring the GlobalProtect Gateway device:

- 1. You should have already generated and imported the gateway server certificate and the portal root CA certificate. Refer to "Certificates and the OCSP Responder" on page 325.
- 2. Create a tunnel interface that is a logical interface used to terminate VPN tunnels. You can use the pre-defined tunnel interface, or create a new one. The interface also needs to be bound to the virtual router and placed in a security zone. It is recommended that you create a new zone, so you can control policies between the branch office and the zone that contains your protected resources. You could put the tunnel interface in the same zone as the protected resources, which will give the satellite

devices access to the resources, but that will not give you granular policy control for the satellite networks.

Refer to "Configuring Tunnel Interfaces" on page 147.

- 3. On the firewall that will be used as a GlobalProtect Gateway, navigate to **Network > GlobalProtect > Gateways** and click the **General** tab.
- 4. Click **Add** and name the gateway.
- 5. Select the interface with Internet connectivity that satellite devices will connect to in order to establish VPN connectivity and also select the **IP Address**.
- 6. Select the gateway server certificate that was signed by the portal and imported to the gateway.
- 7. Select the certificate profile in the **Certificate Profile** drop-down.
- 8. Click the **Satellite Configuration** subtab and check the **Tunnel Configuration** box.
- 9. Select the tunnel interface you created or select the default tunnel interface. You can also set **Replay attack detection** or **Copy TOS** (optional).
- 10. Set the **Configuration Refresh Interval**. This is the setting that each satellite will use to determine how often the satellite device should check the portal for configuration updates, such as the addition of a new gateway that the satellites can use. The default is 24 hours and the range is 1-48 hours.
- 11. Click the **Tunnel Monitoring** check box and enter a **Destination IP**. You can enter any desired IP address that is reachable, or if you leave this field blank, the monitoring profile will ping the gateway's tunnel interface IP address.
- 12. In the **Tunnel Monitor Profile** drop-down, configure a new profile and make sure the action is **failover**.
- 13. In the **IPSec Crypto Profile** you configure a new profile or select the default.
- 14. Click the Network Settings subtab and populate the DNS, IP Pool and Access route information that the satellite devices will use. The IP Pool is a range of IP addresses that will be assigned to the satellite device's tunnel interface that connects to the gateway when a VPN is established. The pool must be large enough to manage all satellites that will connect to the gateway. For the primary and secondary DNS fields, you can select Inheritance Source and select an existing interface that will be used to provide this information to the satellite devices.
- 15. Enter any desired **Access Route** information that you would like to have pushed down to the satellite devices. For example, you could provide access routes to implement split tunneling, so Internet traffic from the branch office does not pass through the VPN. If no access routes are provided, all traffic will pass through the tunnel.
- 16. Click the **Route Filter** tab and either leave this option unchecked to accept all routes sent by the satellite, or click **Accept published routes** and provide a list of subnets, which will cause all other routes to be filtered out. If you are using NAT, which will use the tunnel interface IP, you don't need to filter since no routes need to be published to the gateway.
- 17. Click **Commit** to activate the gateway configuration.

GlobalProtect Portal Configuration

The portal is used to manage the VPN environment. It handles the enrollment process for new satellite devices, maintains the list of gateways to be used by satellites, and can acts as the root CA for the gateways and satellites. You can also use an external certificate authority (CA) if desired.

Once the satellite device receives its initial configuration and certificates from the portal, a VPN connection is established with all defined gateways, so it's not required to have a VPN between the portal and the gateways. The only time the gateway needs to communicate with the portal is for OCSP requests, which only requires HTTP access.

You can configure the portal as a gateway as well. This will allow you to have a gateway access point for satellite devices at the same location as the portal. You can then add additional external gateways for other locations throughout your global VPN network.

- "Create an Authentication Profile" on page 330
- "Security Zone and Interface" on page 331
- "Configuring the GlobalProtect Portal" on page 331

Create an Authentication Profile

The authentication profile determines the authentication method that will be used by satellite devices to connect to the portal. When using this method, the satellite device administrator will enter a login and password when connecting to the portal for the first time. Once the device is authenticated, it will add the serial number and host name to the portal. You can use a Local DB account, RADIUS, LDAP using active directory, or Kerberos. In this case, the portal administrator does not need to add the satellite device serial number in the portal satellite configuration. Refer to "Authentication Profiles" on page 62.

You can also add the serial number of the satellite device as described in "Configuring the GlobalProtect Portal" on page 331. In both cases though, you must have an authentication profile defined in order to commit your configuration. If you manually add a device serial number to the portal configuration, the authentication profile is not used, but it is still required.

Configure the authentication profile:

- 1. To use a local account for authentication, on the portal navigate to **Device > Local Users Database > Users** and click **Add**.
- 2. Enter a username in the **Name** field and then enter a password in the Password/Confirm Password fields. You can also use a password hash by selecting the Password Hash button and pasting a password hash in the field.
- 3. Click **OK** to create the account.
- 4. Navigate to **Device > Authentication Profile** and click **Add**.
- 5. Enter a profile name in the **Name** field and in the **Authentication** drop-down, select **Local Database**. If you will be using RADIUS, LDAP, or Kerberos, select one of those options and configure the server. Once configured, you can then select usernames that are available on those systems.
- 6. In the Allow List, click Add and then select the local user you created.

Security Zone and Interface

On each device that will participate in the VPN (Portal, Gateway, Satellite), you will need to configure an interface that will have access to the other devices in the VPN environment and put them in the desired security zone. This will allow each device to communicate over the WAN. It is recommended that you create a new zone, so you can control policies between the branch office and the zone that contains your protected resources.

The initial connection from the satellite to the portal for the certificate download and initial configuration will use SSL. Once the satellite is configured, it will establish an IPSec VPN with the configured gateways. Make sure that the security zones that you create have the Enable User Identification check box enabled, so users from the branch office can be identified if User-ID is configured. This is required for zones that will be used on the GlobalProtect Gateway. Refer to "Security Zones" on page 151.

Configuring the GlobalProtect Portal

The portal configuration defines the network interface that will be used for satellite enrollment, authentication, and appearance of the portal. It can also act as a root CA and OCSP responder.

- 1. Navigate to Network > GlobalProtect > Portals.
- 2. Click **Add** and name the portal.
- Select the interface that will act as the ingress interface for the satellite devices and then 3. select the Internet accessible IP address from the drop-down.
- 4. Select the portal server certificate that you created earlier. Refer to "Certificates and the OCSP Responder" on page 325.
- Select the Authentication Profile that you created earlier. The authentication profile is 5. required, even though the satellite device will not use it if you add the satellite device serial number to the portal configuration.
- 6. In this case, a client certificate is not required on the portal since the portal automatically generates client certificates for satellite devices using the root CA.
- 7. Select a certificate profile to be used for satellite authentication. Alternatively, you can enter the satellite serial number in the **Satellite Configuration** sub tab, so the satellite administrator does not have to enter login credentials.
- You can optionally configure appearance settings. 8.
- Click the Satellite Configuration tab that will define options for the satellite devices that 9. will connect to the portal.
- 10. Click Add and name the GlobalProtect Satellite profile. This profile will define the satellite devices, enrollment method, and the list of gateways that the satellites will use in the VPN environment.
- 11. Set the **Configuration Refresh Interval**. This is the setting that each satellite will use to determine how often the satellite device should check the portal for configuration updates, such as new gateways being added.
- 12. In the Devices tab, you can manually add satellite device serial numbers, or the list will be updated when a satellite device connects to the portal for the first time using a login and password authentication.

The next step provides information on the enrollment methods.

- 13. The portal can be configured to allow two types of authentication methods for the initial enrollment (you can configure both of the below options):
 - Manually add the satellite device serial number to the portal in the Network > GlobalProtect > Portals > Satellite Configuration page. Click Add > Devices then click Add again and enter the satellite device serial number. When the satellite device is configured at the branch office and establishes the initial connection to the portal, no authentication is required. Once the satellite device connects, the host name will be added to the portal configuration automatically.
 - Use a login name and password prompt from the satellite device by adding an enrollment user or group to the portal in the Network > GlobalProtect > Portals > Satellite Configuration page. Click Add > Enrollment User/User Group then click Add again and select a user or group from the drop down list or select Any. When using this method, you do not need to enter the satellite device information on the portal. When the device connects, the administrator configuring the satellite device will be prompted to login. Once authenticated, the satellites serial number and host name will automatically be added to the portal satellite list.
- 14. In the Gateways tab, enter the IP address or hostname of the gateway(s) that the defined list of satellite devices will use for VPN connectivity. You can also set the routing priority of the gateway from this page. If the satellite device has tunnels to multiple gateways, duplicate routes may exist. By setting a priority, the gateway with the highest priority will be used first. If the gateway fails, the next highest priority will be used. The priority range is from 1-25, since 25 is the limit on the number of gateways that a satellite can use.
- 15. Click **OK** to save your changes and then under **Trusted Root CA**, click **Add** and select the portal root CA.
- 16. In the **Issuing Certificate** drop-down, also select the portal root CA. This is the certificate that will be used to automatically generate certificates for the satellite devices during enrollment.
- 17. You can modify the validity period and certificate renewal period if desired. The validity period defines the lifetime of the certificate and the renewal period defines how often the certificate will be renewed. If you do not want your certificates to expire, just make sure the renewal period is lower than the validity period.
- 18. In the **OCSP Responder** drop-down, select the portal's OCSP responder that you created earlier and then click **OK** to save your changes.
- 19. Click **Commit** to activate your changes.

GlobalProtect Satellite Configuration

On the branch office satellite device, you first configure an interface with WAN connectivity and set up a security zone and policy to allow the branch office LAN to communicate with the Internet. The satellite device is then configured with the host or IP address of the portal. Once the configuration is committed, the satellite will communicate with the GlobalProtect Portal, certificates are signed and downloaded with the initial VPN configuration and routing information. The satellite then connects to all configured GlobalProtect gateways to establish VPN connectivity from the branch to the corporate office. Refer to "GlobalProtect Satellite" on page 317 for field descriptions).

- 1. On the portal devices, you should have already entered the serial number of the satellite device you are configuring, or have created a username and login that can be used for the initial configuration. Refer to "Configuring the GlobalProtect Portal" on page 331.
- 2. Navigate to **Network** > **IPSec Tunnels** page.
- 3. Click Add to create a new IPSec Tunnel.
- 4. Enter a **Name** for the tunnel profile and select the default **Tunnel Interface** from the drop down, or create a new tunnel interface.
- 5. Choose **Type GlobalProtect Satellite** and enter the **Portal Address** (IP or hostname).
- 6. Select the **Interface** to be used to connect to the portal and then select the **Local IP Address** of the interface. In the example illustration shown in Figure 44 on page 324, the **Portal Address** would be "portal.paloaltonetworks.com".
- Click the Advanced tab and configure the desired advanced options. It is recommended that you enable Publish all static and connected routes to Gateway, so that routes are shared between all devices in the VPN network.
 Note: If this branch office has the subnets as other branch offices in the VPN, do not enable this option unless you are using NAT.
- 8. Click **Commit** to activate the changes on the satellite device.

You can now test connectivity between the satellite devices and gateways. To view the status of the VPN on the satellite device, navigate to **Network > IPSec Tunnels**. You should see a green status indicator in the **Status** column.

Dynamic Routing Protocols and Large Scale VPNs

When using dynamic routing protocols over your large scale VPN network, the following behavior should be noted in relation to flooding and replication:

- Routing Information Protocol (RIP)—Control packets are flooded to all tunnels (spokes) on the egress interface (hub to all spokes). Control packets are not replicated to remaining tunnels (spokes) when one tunnel receives a packet from the incoming spoke to the remaining spokes.
- OSPF—It is recommended that you use Point to Multipoint (P2MP). In broadcast mode, the replication behavior is the same as RIP.

Multicast replication behavior is the same as RIP and data packets are not replicated. Also, multicast distribution in a large scale VPN environment is supported where the multicast source resides behind the gateway and receivers reside behind the satellite VPN devices. Deployments in which the multicast source resides in a satellite location are not currently supported.

Backing up a GlobalProtect Portal

In a large scale VPN configuration, the GlobalProtect Portal is a critical piece of the environment that manages the VPN configuration, all certificates used for authentication, and the list of satellite firewalls that are part of the VPN network. The certificate and satellite information is dynamic and changes often and is not part of the device configuration.

The "Export device state" feature is used to export the configuration and dynamic information from a firewall that is configured as a GlobalProtect Portal with the large scale VPN feature enabled. The export file contains a list of all satellite devices managed by the Portal, the running configuration at the time of the export, and all certificate information (Root CA, Server, and Satellite certificates). If the Portal experiences a failure, the export file can be imported to restore the Portal's dynamic information.

Important: You must manually run the device state export from **Device > Setup > Operations** by clicking the **Export device state** option, or you can create a scheduled XML API script to export the file to a remote server. This should be done on a regular basis since satellite certificates may change often.

To create the device state file from the CLI, run save device state. The file will be named device_state_cfg.tgz and is stored in /opt/pancfg/mgmt/device-state. The operational command to export the device state file is scp export device-state (you can also use tftp device-state export).

For information on using the XML API, refer to the document "PAN-OS XML-Based Rest API Usage Guide" at *https://live.paloaltonetworks.com/community/documentation*.

Chapter 9 Configuring GlobalProtect

This chapter describes GlobalProtect, which allows secure login from client systems located anywhere in the world:

- "Overview" in the next section
- "Setting Up GlobalProtect" on page 337
- "Setting Up and Activating the GlobalProtect Agent" on page 352

Overview

GlobalProtect provides security for client systems, such as laptops, that are used in the field by allowing easy and secure login from anywhere in the world. With GlobalProtect, users are protected against threats even when they are not on the enterprise network by sending their traffic through a Palo Alto Networks firewall that is within close geographic proximity. The user's access level is determined by a host information profile (HIP) that notifies the firewall about the user's local configuration. The HIP information can be used for granular access control based on the security programs that are running on the host, registry values, and many other checks such as whether the host has disk encryption enabled.

The GlobalProtect agent can also be a Palo Alto Networks satellite device (firewall), but instead of agent software being downloaded, only the necessary certificate for authentication and VPN configuration is needed. For information on satellite devices, refer to "GlobalProtect Large Scale VPN Deployment" on page 323.

The following elements are used to provide GlobalProtect functionality:

- **Portal**—Palo Alto Networks firewall that provides centralized management for the GlobalProtect system.
- **Gateways**—Palo Alto Networks firewalls that provide security enforcement for traffic from GlobalProtect agents.
- Agent—Application that is installed on the client system and configured to connect to the portal and gateways to provide network access for user systems. The agent also provides information about the user's local configuration to the gateways.
- **Satellite**—A satellite device is a Palo Alto Networks firewall that is typically installed at a branch office and is initialized and managed by a GlobalProtect portal. After the satellite receives its initial configuration from the portal, it connects to all configured gateways to

establish VPN connectivity. By using GlobalProtect, firewalls can be quickly installed at branch offices with very little initial configuration. Refer to "GlobalProtect Large Scale VPN Deployment" on page 323.

The connection process works as follows:

- 1. The user downloads the GlobalProtect agent from the portal onto a client system, such as a laptop. The agent connects to the portal through an SSL connection and downloads a configuration file. This configuration file includes information about the various GlobalProtect gateways that can be used for connectivity.
- 2. The agent performs a reverse Domain Name System (DNS) lookup to determine whether the client system is currently on the internal enterprise network or on an external network.
- 3. If the connection is to the external network, the agent attempts to make SSL connections to all external gateways and then selects the best gateway.
- 4. Based on the configuration, an IPSec or SSL tunnel is established between the agent and the gateway and a default route is inserted to direct all traffic through the tunnel for the purpose of policy control and threat scanning.
- 5. The agent submits the client system posture, the HIP.

GlobalProtect allows the use of HIP profiles in security profiles. A HIP object specifies a set of criteria for the client system that is treated as a unit when defining HIP profiles. For example, a HIP object might specify full disk encryption or software patching.

HIP profiles can incorporate HIP objects on a match or no-match basis. For example, a HIP profile might include a match for HIP objects that specify that the client system includes both full disk encryption and certain software patches to be installed.

The portal stores client configurations and maintains internal and external gateways lists. It also lists the CA that is used by gateways and agents for mutual authentication.

Each gateway can operate in tunnel mode (external gateways) or non-tunnel mode (internal gateways). Gateways in non-tunnel mode receive only the HIP from clients. All communication is between agents and gateways; there is no inter-gateway communication.

Policies are enforced on gateways based on user or HIP information, and HIP objects and profile matches are logged in the gateway's HIP database. This information is displayed in the Application Control Center (ACC), logs, and custom reports.

GlobalProtect Authentication

Connectivity between all parts of the GlobalProtect infrastructure is authenticated using SSL certificates. The portal can act as a certificate authority (CA) for the system (using a self-signed or imported subordinate issuing CA certificate within the portal), or customers can generate certificates using their own CAs. It is recommended (but not required) that the portal, gateways, and the GlobalProtect software agents use certificates signed by the same CA. Prior to transferring any information, the agent verifies the portal server certificate. If the certificate presented by the portal is not signed by a Trusted CA, the agent shows a warning dialog and waits for the user response to continue or cancel.

As part of the configuration bundle that is sent to the client, the portal includes the public certificate of the CA and the needed client certificate and key. The client certificate is used by GlobalProtect gateways to authenticate and identify the client.

If an internal CA is used, the certificate is auto-generated and does not require user interaction. The portal can export the necessary server certificate and key for the gateways. If an external CA is used, support is provided to import the CA certificate along with a server certificate and key for the portal and gateways and a client certificate and key for the clients.

For more information on authentication, refer to "Authentication Profiles" on page 62 and "Authentication Sequence" on page 67.

Setting Up GlobalProtect

Setting up GlobalProtect on the firewall involves the following tasks:

- 1. Define HIP objects, as described in "Setting Up HIP Objects" on page 337.
- 2. Create HIP profiles, as described in "Setting Up HIP Profiles" on page 340.
- 3. Set up the portal, as described in "Setting Up the GlobalProtect Portal" on page 341.
- 4. Set up gateways, as described in "Setting Up the GlobalProtect Gateways" on page 347.
- 5. Define security policies that include HIP profiles, as described in "Defining Security Policies" on page 187.
- 6. Distribute the GlobalProtect agent, as described in "Setting Up the GlobalProtect Agent" on page 353.
- 7. Monitor client activity, as described in "Viewing the Logs" on page 264.

Setting Up HIP Objects

Objects > GlobalProtect > HIP Objects

Use this page to define settings for use in HIP profiles for GlobalProtect. Each HIP object defines a set of criteria for the client system that is treated as a unit when defining HIP profiles.

Field	Description
General Tab	
Name	Enter a name for the HIP object (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Shared	Select the check box to make the object available to all virtual systems.
Description	Enter an optional description.
Host Info	Select the check box to specify host information.
Domain	To match domains, choose an operator from the drop-down list and enter a string to match.
OS	Use the drop-down lists to specify the operating system (OS) for the GlobalProtect agent.
Client Versions	To match OS versions on the client, choose an operator from the drop- down list and enter a string to match.

Table 137. HIP Object Settings

Field	Description
Patch Management Tab	
Patch Management	Select the check box to include software patch management in the HIP. When the check box is selected, the settings are activated. You can then define patch management vendors, such as Microsoft, and only allow clients that have certain patch levels installed.
Criteria	Specify the following settings on this subtab:
	• Is Enabled —Choose whether the settings on this tab are enabled (yes) or disabled (no), or not available.
	• Is Installed —Select the check box to verify if the client has the defined software patch installed.
	• Severity—Choose the level of importance for missing patches.
	 Check—Choose how the system should check for patches.
	• Patches—Click Add and enter patch file names.
Vendor	Click Add to specify patch management products. Choose a vendor from the drop-down list, and then click Add to choose a specific product. Click OK to save the settings and return to the Patch Management tab.
Firewall Tab	
Firewall	Select the check box to activate this tab and then specify the following settings:
	• Is Enabled —Choose whether the settings on this tab are enabled (yes) or disabled (no), or not available.
	• Is Installed—Select the check box to verify if the client has the defined firewall software installed.
	• Vendor and Product—Click Add to specify specific firewalls. Choose a vendor from the drop-down list, and then click Add to choose a specific firewall version. Click OK to save the settings and return to the Firewall tab.
	• Exclude Vendor—Select the check box if you want to exclude rather than include the specified vendors and products.

Table 137. HIP Object Settings (Continued)

Field	Description
Antivirus Tab	
Antivirus	Select the check box to activate this tab and then specify the following settings:
	• Real Time Protection —Choose whether to require real-time protection (if available).
	• Is Installed—Select the check box to verify if the client has the defined antivirus software installed.
	• Virus Definition Version—Choose from the drop-down list. If you choose Within or Not Within, specify the number of Days or Hours to match.
	 Product Version—Choose an operator from the drop-down list and specify a matching string.
	• Last Scan Time—Choose from the drop-down list. If you choose Within or Not Within, specify the number of Days or Hours to match.
	• Vendor and Product—Click Add to specify antivirus products. Choose a vendor from the drop-down list, and then click Add to choose a specific product. Click OK to save the settings and return to the Antivirus tab.
	• Exclude Vendor—Select the check box if you want to exclude rather than include the specified vendors and products.
Anti-spyware Tab	
Anti-spyware	Select the check box to activate this tab and then specify the following settings:
	• Real Time Protection —Choose whether to require real-time protection (if available).
	• Is Installed—Select the check box to verify if the client has the defined anti-spyware software installed.
	• Virus Definition Version—Choose from the drop-down list. If you choose Within or Not Within, specify the number of Days or Hours to match.
	 Product Version—Choose an operator from the drop-down list and specify a matching string.
	• Last Scan Time—Choose from the drop-down list. If you choose Within or Not Within, specify the number of Days or Hours to match.
	• Vendor and Product—Click Add to specify anti-spyware products. Choose a vendor from the drop-down list, and then click Add to choose a specific product. Click OK to save the settings and return to the Anti- spyware tab.
	• Exclude Vendor—Select the check box if you want to exclude rather than include the specified vendors and products.

Table 137. HIP Object Settings (Continued)

Field	Description
Disk Backup Tab	
Disk Backup	Select the check box to activate this tab and then specify the following settings:
	• Is Installed—Select the check box to verify if the client has the defined backup software installed.
	• Last Backup Time—Choose from the drop-down list. If you choose Within or Not Within, specify the number of Days or Hours to match.
	• Vendor and Product—Click Add to specify disk backup products. Choose a vendor from the drop-down list, and then click Add to choose a specific product. Click OK to save the settings and return to the Disk Backup tab.
	• Exclude Vendor—Select the check box if you want to exclude rather than include the specified vendors and products.
Disk Encryption Tab	
Disk Encryption	Select the check box to activate this tab and then specify the following settings:
Criteria	Specify the following settings on this subtab:
	• Is Installed —Select the check box to verify if the client has the defined disk encryption software installed.
	• Encrypted Locations—Click Add to specify the drive or path that refers to an encrypted data store:
	- Encrypted Locations—Choose the location from the drop-down list.
	 State—Specify the state of the encrypted location by choosing an operator and value from the drop-down list.
	Click OK to save the settings and return to the Disk Encryption tab.
Vendor	Click Add to specify specific disk encryption products. Choose a vendor from the drop-down list, and then click Add to choose a specific product. Click OK to save the settings and return to the Disk Encryption tab.
Custom Checks Tab	
Process List	Click Add to specify the list of processes to be checked on the client system to see if they are running. For example, to determine whether a software application is running, add the name of the executable file to the process list.
Registry Key	Click Add to specify that a particular registry key is present or has a specified value.
Plist	Plists are preferences files on MacOS. Define the path to a specific plist file. You can also include preference keys and values to check for within the file.

Table 137. HIP Object Settings (Continued)

Setting Up HIP Profiles

► Objects > GlobalProtect > HIP Profiles

After defining HIP objects (refer to "Setting Up HIP Objects" on page 337), use this page to create HIP profiles for GlobalProtect. When defining HIP profiles, you specify match criteria that are built from the previously-defined HIP objects.

Field	Description
Name	Enter a name for the profile (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Description	Enter an optional description.
Shared	Select the check box to make the profile available to all virtual systems.
Match	Define one or more HIP objects that you want to check for on the client. Enter the HIP objects to include, or click Add Match Criteria to create objects. When including multiple HIP objects, you can use AND, OR and NOT operators to create a Boolean expression. Using this method, you can establish complex HIP profiles, for example, to test if your clients have antivirus installed AND disk encryption installed and enabled.

Table 138. HIP Profile Settings

Setting Up the GlobalProtect Portal

Network > GlobalProtect > Portals

Use this page to configure portals for GlobalProtect.

Table 139. GlobalProtect Portal Settings

Field	Description
Portal Configuration to	ıb
Name	Enter a name for the portal (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Location	Select the virtual system, if the multiple virtual systems option is enabled.
Network Settings	
Interface	Select the firewall interface that will be used as the ingress for remote clients/firewalls.
IP Address	Specify the IP address on which GlobalProtect portal web service will be running.
Server Certificate	Select the SSL certificate for the GlobalProtect portal.
Authentication	
Authentication Profile	Choose an authentication profile to authenticate access to the portal. Refer to "Authentication Profiles" on page 62.
Client Certificate	Select the certificate the client will use to connect to the gateways.
Certificate Profile	Select the certificate profile that is used to authenticate smartcard users on the portal.
Appearance	
Custom Login Page	Choose an optional custom login page for user access to the portal.
Custom Help Page	Choose an optional custom help page to assist the user with access to the portal.

Field	Description
Client Configuration tab	

Table 139. GlobalProtect Portal Settings (Continued)

Field	Description
General subtab settings	Click Add to display the subtabs, and specify the following settings on the General subtab:
	• Name—Enter a name to identify this client configuration.
	• Use single sign-on—Select the check box to have GlobalProtect use the users' Windows login credentials to transparently connect and authenticate to the GlobalProtect portal and gateways. No username and password is needed in the GlobalProtect agent configuration.
	• Config Refresh Interval (hours) —Specify the interval in days to refresh the GlobalProtect agent configuration (default 24 hours, range 1-168days).
	Connect Method:
	 on-demand—Select this option to allow users to establish a connection on demand. With this option, the user must explicitly initiate the connection. This function is primarily used for remote access connections.
	 user-logon—When this option is set, the GlobalProtect agent will automatically establish a connection after users log in to their computers. If you select Use single sign-on, the username and password used to log in to Windows is captured by the GlobalProtect agent and used to authenticate.
	 pre-logon—Select this option to preserve pre-logon and post-logon services provided by a corporate infrastructure regardless of where the user machine is located. GlobalProtect will establish a connection prior to user login to the computer. By doing this, a company can create a "logical network" that maintains the security and management features normally achieved by a physical network. Tunnel selection and establishment happens pre-logon based on machine certificates that need to be pre-deployed on client systems outside of GlobalProtect.
	Active Directory group policy enforcement, maintaining drive mapping to server resources, and the ability to receive central software deployment downloads while remote. One specific example of how the pre-logon feature works is if a remote user forgets his/her password, since GlobalProtect would connect and use the cached credentials and establish a VPN before the login prompt even appears, a domain administrator could reset the user's password as if they were logged in directly to a domain controller on the physical network. You could also push new Active Directory Group Policies
	With the pre-logon option enabled, you can also use this feature in from the Users tab when configuring Security Policies or in the Source tab in Decryption policies. For example, you may want to set a security policy to only allow access to AD servers and DNS servers on your network for machines that are in the pre-logon state. Note that each user is identified with the name pre-logon and you would need to check their machine name in HIP reports to determine specific users.
	 Third Party VPN—Click Add to add a list of third party remote access VPN clients that might be present on the system. If configured, GlobalProtect will ignore those clients and their route settings to ensure that it does not interfere or conflict with them.

Table 139. GlobalProtect Portal Settings (Continued)

Field	Description
General subtab (Continued)	 Internal Host Detection—With this option, GlobalProtect tries to resolve the configured hostname to the configured IP address. If this fails, GlobalProtect assumes the computer to be outside of the corporate network and will establish a tunnel with any of the available external gateways configured in the Gateways tab. Select the check box to enable internal host detection using DNS lookup. Specify the following:
	 IP Address—Enter an internal IP address for the internal host detection.
	 Hostname—Enter the hostname that resolves to the above IP address within the internal network.
User/User Group subtab settings	Specify the user or user group to which the particular client configuration is applied.
Gateways subtab	Specify gateway settings:
	• Cutoff Time —Specify the timeout (seconds) after which the agent will dismiss a gateway response. The agent dismisses gateway responses after either the configured cutoff time or the socket timeout is reached. If 0 is specified, the cutoff time is ignored by the agent.
	 Internal Gateways—Specify the internal firewalls that the agent will authenticate and provide HIP reports to.
	• External Gateways—Specify the list of firewalls the agent will try to establish a tunnel with when not on the corporate network. The agent will contact all of the gateways and establish a tunnel with the firewall that provides the fastest response and the lowest priority value. Select the Manual check box if you want to allow users to manually select a different gateway while their VPN is connected. The GlobalProtect agent will have the option to connect to any external gateway that has the manual selection. When connecting to the new gateway, the existing tunnel will be disconnected and a new tunnel will be established. The manual gateways can also have different authentications mechanism than the primary gateway. If the client system is restarted, or if a rediscovery is performed, the GlobalProtect agent will connect to the primary gateway. This feature is useful if you have a group of users who need to temporarily connect to a specific gateway to access a secure segment of your network.
	Note: You can only connect to gateways defined in the portal, you cannot connect to different portals and gateways.

Table 139. GlobalProtect Portal Settings (Continued)

Field	Description
Agent subtab	Specify the following settings:
	• Enable advanced view—Deselect this check box to restrict the user interface on the client side to the basic minimum view. By default, the advanced view setting is enabled on all GlobalProtect agents.
	 User can save password—Select the check box to allow users to save their passwords.
	 Agent User Override—Select an override option:
	 disabled—User override is disabled.
	 with-comment—The user is prompted to enter a comment when disabling the GlobalProtect agent.
	 with-passcode—The user must provide the passcode to use the GlobalProtect agent override.
	 with-ticket—This option enables a challenge-response mechanism to authorize disabling GlobalProtect on the client side. When this option is selected, the user is prompted with a challenge when disabling GlobalProtect. The challenge is then communicated to the firewall administrator out-of-band, and the administrator can validate the challenge through the firewall management interface. The firewall produces a response that is read back to the user who can then disable GlobalProtect by entering the response in GlobalProtect.
	• Agent User Override Timeout—Specify the maximum wait time (sec- onds) before data collection times out.
	• Max Agent User Overrides—Specify the maximum number of times a user can disable GlobalProtect before a successful connection to a fire-wall is required.
	• Display welcome page —Select the check box to allow display of a welcome page for the portal.
	• Welcome Page—Choose the factory default welcome page, or click Import to import another page. If you choose None and select the Display Welcome Page option, a blank page is displayed.
	• Enable rediscover network option—Select this check box to allow the user to manually trigger network rediscovery.
	• Enable resubmit host information option—Select this check box to allow the user to manually trigger resubmission of the latest HIP.
	• Client Upgrade —Specify whether to prompt the client to update after configuration changes (prompt) or to perform the upgrade silently without a user prompt (transparent).

Table 139. GlobalProtect Portal Settings (Continued)

Field	Description
Data Collection Subtab	Specify the following settings on this subtab:
	• Max Wait Time—Specify the maximum wait time (seconds) before data collection times out.
	• Exclude Categories—Click Add to specify particular software and client configuration categories to exclude from the data collection. Choose a vendor from the drop down list and click Add to choose a specific product. Click OK to save settings.
	• Custom Checks—Specify the following information:
	 Registry Key—(Windows) Click Add to specify that a particular registry key is present or has a specified value.
	 Plist—(Mac) Click Add to specify that a particular plist key is present or has a specified value.
	 Process List—Click Add to specify the list of processes to be checked on the end user systems to see if they are running. For example, to determine whether a software application is running, add the name of the executable file to the process list.
Trusted Root CA	Specify the Root CA or issuing certificates that the GlobalProtect agent will trust when connecting to a gateway. If a gateway presents a certificate to the client that hasn't been issued by one of the listed CAs, the client will reject the handshake and terminate the connection.
	Click Add to specify a root CA certificate.
Agent User Override Key/Confirm Agent User Override Key	When you create a GlobalProtect Agent override with ticket, which can be used to allow users to disable the Agent, you can use this hash key to validate the client when they try to disable with the ticket.
Satellite Configuration	For information on using GlobalProtect for deploying large scale VPNs, refer to "GlobalProtect Large Scale VPN Deployment" on page 323.
General subtab	Click Add to display the subtabs, and specify the following on the GlobalProtect Satellite > General subtab:
	• Name—Enter a name to identify the GlobalProtect satellite device pro- file.
	• Configuration Refresh Interval (hours) —Specify how often satellite devices should check the portal for configuration updates (default 24 hours, range 1-48 hours).
Devices subtab	Click Add to manually add a satellite device using the device serial number. If you use this option, when the satellite device first connects to receive the authentication certificate and the initial configuration, no user login prompt is required. After the satellite device authenticates, the Name (host name) will be added automatically to the Portal.
Enrollment User/User Group subtab	Click Add to add a User or User Group that will be used to authenticate the satellite device when it first connects to the portal. Select Any to allow authentication by any user or user group. When the satellite device first connects to the portal, the administrator will be prompted for a login/ password.

Table 139. GlobalProtect Portal Settings (Continued)

Field	Description
Gateways subtab	Click Add to enter the IP address or hostname of the gateway(s) that will be part of the VPN network. Each satellite device that is added to the portal will create an IPSec tunnel to each gateway. Routing information is also shared between the satellites and gateways.
	You can also define a routing priority to gateways. When a satellite has VPN tunnels to multiple gateways, duplicate routes may exist. By setting a priority, the gateway with the highest priority will be used first. If the gateway fails, the next highest priority will be used. The priority range is from 1-25, since 25 is the limit on the number of gateways that a satellite can connect to.
	The satellite will also share its network and routing information with the gateways if the Publish all static and connect routers to Gateway (configured on the satellite in Network > IPSec tunnels > GlobalProtect Satellite > Advanced tab) option is selected. Refer to "GlobalProtect Satellite" on page 317.
	When adding a gateway, in the portal configuration, you create a new key and certificate using the same issuing CA used to create satellite and other gateway certificates. You export the certificate as a PKCS12 file and import it into the gateway together with the issuing CA certificate.
Trusted Root CA	Click Add and select an existing Trusted Root CA (Certificate Authority), or generate a new Trusted Root CA that will be used to issue client certificates to all satellite and gateway devices for the GlobalProtect satellite profile.
Issuing Certificate	Defines the certificate that the portal will use to issue new IPSec client certificates to satellite and gateway devices.
Validity Period (days)	Specify the issued GlobalProtect satellite certificate lifetime (default 7 days, range 7-365 days).
Certificate Renewal Period (days)	Specify the GlobalProtect satellite certificate renewal period (default 3 days, range 3-30 days). This will determines how often certificates should be renewed.
OCSP Responder	Select the OCSP responder that will be used by the portal to determine the revocation status of the gateway and satellite devices.

Table 139. GlobalProtect Portal Settings (Continued)

Setting Up the GlobalProtect Gateways

Network > GlobalProtect > Gateways

Use this page to configure gateways for GlobalProtect. The gateway can be used to provide VPN connections for client PCs or GlobalProtect satellite devices.

Field	Description
General tab	
Name	Enter a name for the gateway (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Location	Select the virtual system, if the multiple virtual systems option is enabled.
Network Settings	

Table 140. GlobalProtect Gateway Settings

Field	Description
Interface	Select the firewall interface that will be used as the ingress for remote clients/firewalls.
IP Address	Specify the IP address on which GlobalProtect portal web service will be running.
Server Certificate	Choose the server certificate for the gateway.
Authentication	
Authentication Profile	Choose an authentication profile or sequence to authenticate access to the portal. Refer to "Authentication Profiles" on page 62.
Certificate Profile	Choose the certificate profile for client authentication.
Client Configuration tab	
Tunnel Settings subtab	
Tunnel Mode	Select the check box to enable tunnel mode and specify the following settings:
	• Tunnel Interface —Choose the tunnel interface for access to the gateway.
	• Max Users—Specify the maximum number of users that can access the gateway at the same time. If the maximum number of users is reached, subsequent users are denied access with an error message indicating that the maximum number of users has been reached.
	• Enable IPSec—Select the check box to enable IPSec mode for client traffic, making IPSec the primary and SSL-VPN the fall back method.
	• Enable X-Auth Support—Select the check box to enable Extended Authentication (X-Auth) support in the GlobalProtect gateway when IPSec is enabled. With X-Auth support, third party IPSec VPN clients that support X-Auth (such as the IPSec VPN client on Apple iOS and Android devices) can establish a VPN tunnel with the GlobalProtect gateway. The X-Auth option simply provides remote access from the VPN client to a specific GlobalProtect gateway and does not provide the full control and HIP features that are part of the GlobalProtect agent installed on PCs. For details on configuring a GlobalProtect gateway to allow IPSec X-Auth clients, refer to the tech notes "GlobalProtect Config for Apple iOS VPN" and "GlobalProtect Config for Android VPN" in the Tech Note area at https://live.paloaltonetworks.com/community/ documentation.
	 If the group name and group password are specified, the first auther tication phase requires both parties to use this credential to authen- ticate. The second phase requires a valid user name and password, which is verified through the authentication profile configured in the Authentication section.
	 If no group name and group password are defined, the first authent cation phase is based on a valid certificate presented by the third party VPN client. This certificate is then validated through the certificate profile configured in the authentication section.

Table 140. GlobalProtect Gateway Settings (Continued)

Field	Description
Timeout Configuration	Specify the following timeout settings:
	 Login Lifetime—Specify the number of days, hours, or minutes allowed for a single gateway login session.
	• Inactivity Logout —Specify the number of days, hours, or minutes after which an inactive session is automatically logged out.
Network Settings subtab	
Inheritance Source	Select a source to propagate DNS server and other settings from the selected DHCP client or PPPoE client interface into the GlobalProtect agents' configuration. With this setting all client network configuration, such as DNS servers and WINS servers, are inherited from the configuration of the interface selected in the Inheritance Source.
Primary DNS Secondary DNS	Enter the IP addresses of the primary and secondary servers that provide DNS to the clients.
Primary WINS Secondary WINS	Enter the IP addresses of the primary and secondary servers that provide Windows Internet Naming Service (WINS) to the clients.
Check inheritance source status	Click the link to see the server settings that are currently assigned to the client interfaces.
	Click Add to enter a suffix that the client should use locally when an unqualified hostname is entered that it cannot resolve. You can enter multiple suffixes by separating them with commas.
Inherit DNS Suffixes	Select this check box to inherit the DNS suffixes from the inheritance source.
IP Pool	Click Add to specify IP pool settings.
	Use this section to create a range of IP addresses to assign to remote users. When the tunnel is established, an interface is created on the remote user's computer with an address in this range.
	Note: The IP pool must be large enough to support all concurrent connections. IP address assignment is dynamic and not retained after the user disconnects. Configuring multiple ranges from different subnets will allow the system to offer clients an IP address that does not conflict with other interfaces on the client.
	The servers/routers in the networks must route the traffic for this IP pool to the firewall.
	For example, for the 192.168.0.0/16 network, a remote user may be assigned the address 192.168.0.10.
Access Route	Click Add to specify access route options.
	Use this section to add routes that will be pushed to the remote user's computer and therefore determine what the user's computer will send through the VPN connection.
	For example, you can set up split tunneling to allow remote users to access the Internet without going through the VPN tunnel.
	If no route is added, then every request is routed through the tunnel (no split tunneling). In this case, each Internet request passes through the firewall and then out to the network. This method can prevent the possibility of an external party accessing the user's computer and then gaining access to the internal network (with the user's computer acting as bridge).

Table 140. GlobalProtect Gateway Settings (Continued)

Field	Description
HIP Notification subtab	
HIP Notification	Click Add to specify notification options. Select Enable to enable the option to match or not match a message.
	Choose a notification option from the Show Notification As section and choose the radio button for a System Tray Balloon or Pop Up Message , and then specify a message to match or not match. Use these settings to notify the end user about the state of the machine, for example, to provide a warning message.
	Note: The HIP notification pages can be formatted in rich HTML, which can include links to external web sites and resource. Use the link icon and in the rich text settings toolbar to add links.
Satellite Configuration tab	
Tunnel Settings subtab	
Tunnel Configuration	Select the Tunnel Configuration check box and select an existing tunnel interface, or click New Tunnel Interface. For information on creating a tunnel interface, refer to "Configuring Tunnel Interfaces" on page 147.
	Replay attack detection—Protect against replay attacks.
	Copy TOS —Copy the (Type of Service) TOS header from the inner IP header to the outer IP header of the encapsulated packets in order to preserver the original TOS information.
	Configuration refresh interval (hours) —Specify how often satellite devices should check the portal for configuration updates (default 24 hours, range 148 hours).
Tunnel Monitoring	Select the Tunnel Monitoring check box to enable monitoring of the VPN tunnels between the satellite device and gateways. This is recommended to allow failover if one of the gateways cannot communicate with the satellite.
	Destination IP —Specify an IP address on the other side of the tunnel that the tunnel monitor will use to determine if the tunnel is working properly.
	Tunnel Monitor Profile —Select the default monitor profile, or create a new profile. A monitor profile is used to perform automatic failover to another tunnel. Refer to "Defining Monitor Profiles" on page 177.
Crypto Profiles	Select an IPSec Crypto Profile , or create a new profile. This will determine the protocols and algorithms for identification, authentication, and encryption for the VPN tunnels. Refer to "Defining IPSec Crypto Profiles" on page 319.
Network Settings subtab	
Inheritance Source	Select a source to propagate DNS server and other settings from the selected DHCP client or PPPoE client interface into the GlobalProtect satellite configuration. With this setting all network configuration, such as DNS servers, are inherited from the configuration of the interface selected in the Inheritance Source.

Table 140. GlobalProtect Gateway Settings (Continued)

Field	Description
Primary DNS Secondary DNS	Enter the IP addresses of the primary and secondary servers that provide DNS to the satellites.
DNS Suffix	Click Add to enter a suffix that the satellite should use locally when an unqualified hostname is entered that it cannot resolve. You can enter multiple suffixes by separating them with commas.
Inherit DNS Suffix	Select this check box to send the DNS suffix to the satellite devices to use locally when an unqualified hostname is entered that it cannot resolve.
IP Pool	Click Add to specify IP pool settings.
	Use this section to create a range of IP addresses to assign to satellite devices. When the tunnel is established, an interface is created on the satellite device with an address in this range.
	Note: The IP pool must be large enough to support all concurrent connections. IP address assignment is dynamic and not retained after the satellite disconnects. Configuring multiple ranges from different subnets will allow the system to offer satellites an IP address that does not conflict with other interfaces on the device.
	The servers/routers in the networks must route the traffic for this IP pool to the firewall.
	For example, for the 192.168.0.0/16 network, a satellite may be assigned the address 192.168.0.10.
Access Route	Click Add to specify access route options.
	Use this section to add routes that will be pushed to the satellites and therefore determining what routes the satellite will use to send traffic through the VPN connection.
	For example, you can set up split tunneling to allow satellites to access the Internet without going through the VPN tunnel.
	If no route is added, then every request is routed through the tunnel (no split tunneling).
	Click Add to enter a route.
Route Filter subtab	Click the Accept published routes check box and enter the satellite subnets that will be added to the gateway's route table when routing information is shared. Satellite subnets that are not part of the list will be filtered out.

Table 140. GlobalProtect Gateway Settings (Continued)

Setting Up and Activating the GlobalProtect Agent

Device > GlobalProtect Client

The **GlobalProtect Client** page lists the available GlobalProtect releases. When the client connects, the system checks the version and installs the currently activated version of the agent software if it is different from the version that is on the client.



Note: For initial download and installation of the GlobalProtect agent, the user on the client system must be logged in with administrator rights. For subsequent upgrades, administrator rights are not required.

To download and activate the GlobalProtect agent:

- 1. Click the **Download** link for the desired release. The download starts and a pop-up window opens to display the progress of the download. When the download is complete, click **Close**.
- 2. To activate a downloaded release, click the **Activate** link for the release. If an existing version of the client software has already been downloaded and activated, a pop-up message is displayed to indicate that the new version will be downloaded the next time that the clients connect.
- 3. To activate the client that was previously installed by way of the **Upload** button, click the **Activate From File** button. A pop-up window opens. Select the file from the drop-down list and click **OK**.
- 4. To remove a downloaded release of the client software from the firewall, click the **Remove** icon in the rightmost column.

Setting Up the GlobalProtect Agent

The GlobalProtect agent (PanGP Agent) is an application that is installed on the client system (typically a laptop) to support GlobalProtect connections with portals and gateways and is supported by the GlobalProtect service (PanGP Service).



Note: Make sure that you choose the correct installation option for your host operating system (32-bit or 64-bit). If installing on a 64-bit host, use 64-bit browser/Java combo for the initial installation.

To install the agent, open the installer file and follow the on-screen instructions. To configure the agent:

1. Choose Start > All Programs > Palo Alto Networks > GlobalProtect > GlobalProtect.

🛃 GlobalProtec	t	
Status Details S	ettings Host State Troubleshooting	
Settings		
Username:		
Password:		
	Remember Me	
Portal:		
	Apply Clear	
Cogin Messages Oct 14 06:13:03) 2010 - Portal Error	

The client interface opens to show the **Settings** tab.

Figure 45. GlobalProtect - Settings Tab

- 2. Specify the username and password to use for GlobalProtect authentication, and optionally select the **Remember Me** check box.
- 3. Enter the IP address of the firewall that serves as the GlobalProtect Portal.
- 4. Click **Apply**.

Using the GlobalProtect Agent

The tabs in the GlobalProtect agent contain useful information about status and settings, and provide information to assist in troubleshooting connection issues.

- Status tab—Displays current connection status and lists any warnings or errors.
- **Details tab**—Displays information about the current connection, including portal IP addresses and protocol, and presents byte and packet statistics about the network connection.
- **Host State tab**—Displays the information stored in the HIP. Click a category on the left side of the window to display the configured information for that category on the right side of the window.

- **Troubleshooting tab**—Displays information to assist in troubleshooting.
 - **Network Configurations**—Displays the current client system configuration.
 - **Routing Table**—Displays information on how the GlobalProtect connection is currently routed.
 - **Sockets**—Displays socket information for the current active connections.
 - Logs—Allows you to display logs for the GlobalProtect agent (PanGP Agent) and service (PanGP Service). Choose the log type and debugging level. Click Start to begin logging and Stop to terminate logging.

Chapter 10 Configuring Quality of Service

This chapter describes how to configure quality of service (QoS) on the firewall:

- "Firewall Support for QoS" in the next section
- "Defining QoS Profiles" on page 358
- "Defining QoS Policies" on page 359
- "Displaying QoS Statistics" on page 362

Firewall Support for QoS

The firewall supports fine grained QoS settings for clear text and tunneled traffic upon egress from the firewall. (Ingress QoS processing is not supported.) QoS profiles are attached to physical interfaces to specify how traffic classes map to bandwidth (guaranteed, maximum) and priority. QoS policy is then used to map specific sessions to QoS classes. QoS classification is supported with all interface types except Aggregate Ethernet.

The firewall supports the following QoS settings:

- Use the Network > QoS page to configure QoS settings for firewall interfaces and specify criteria for the clear text and tunneled traffic that leaves the firewall through those interfaces. Refer to "Configuring QoS for Firewall Interfaces" on page 356.
- For each interface, you can define QoS profiles that determine how the QoS traffic classes are treated. You can set overall limits on bandwidth regardless of class and also set limits for individual classes. You can also assign priorities to different classes. Priorities determine how traffic is treated when contention occurs. Refer to "Defining QoS Profiles" on page 358.
- Use the Policies > QoS page to configure policies to activate the QoS restrictions. Refer to "Defining QoS Policies" on page 359.

Important items to consider when configuring firewall support for QoS

- When setting up the QoS profile, the guaranteed and maximum egress settings defined for the classes must not exceed the guaranteed and maximum egress settings defined for the profile itself.
- Traffic that does not match QoS policy will be assigned a default class of 4. Be sure to assign a maximum guaranteed bandwidth and priority with this in mind.
- Each firewall model supports a maximum number of ports that can be configured with QoS. Refer to the spec sheet for your firewall model at *http://www.paloaltonetworks.com*.

Configuring QoS for Firewall Interfaces

► Network > QoS

Use the **QoS** page to configure bandwidth limits for firewall interfaces.

Field	Description
Physical Interface	
Interface Name	Select the firewall interface.
Egress Max (Mbps)	Enter the limit on traffic leaving the firewall through this interface (Mbps).
Turn on QoS feature on this interface	Select the check box to enable QoS features.
Default Profile:	Select the default QoS profiles for clear text and for tunneled traffic. You
Clear Text	must specify a default profile for each. For clear text traffic, the default
Tunnel Interface	the default profile is applied individually to each tunnel that does not
	have a specific profile assignment in the detailed configuration section. For instructions on defining QoS profiles, refer to "Defining QoS Profiles" on page 358.
Clear Text Traffic and Tunneled Traffic	Specify the following settings on the Tunneled and Clear Text Traffic tabs. These values apply unless they are overridden by setting in the Detail Configuration area, as described later in this table.
Egress Guaranteed (Mbps)	Enter the bandwidth that is guaranteed for tunneled traffic from this interface.
Egress Max (Mbps)	Enter the limit on traffic leaving the firewall through this interface (Mbps).

Table 141. QoS Settings

Field	Description
Add	Click Add from the Clear Text Traffic or Tunneled Traffic tabs to define additional granularity to the treatment of clear text traffic or to override the default profile assignment for specific tunnels. If this section is left blank, the values specified in Group Configuration are used.
	For example, assume a configuration with two sites, one of which has a 45 Mbps connection and the other a T1 connection to the firewall. You can apply restrictive QoS settings to the T1 site so that the connection is not overloaded while also allowing more flexible settings for the site with the 45 Mbps connection.
	To add granularity for clear text traffic, click the Clear Text tab, click Add , and then click individual entries to configure the following settings:
	• Name—Enter a name to identify these settings.
	• QoS Profile —Select the QoS profile to apply to the specified interface and subnet. For instructions on defining QoS profiles, refer to "Defining QoS Profiles" on page 358.
	Note: The QoS rules for clear text are applied in the specified order. Refer to the following tech note for more details on QoS https://live.paloaltonetworks.com/docs/DOC-3439.
	• Source Interface —Select the firewall interface.
	• Source Subnet —Select a subnet to restrict the settings to traffic coming from that source, or keep the default any to apply the settings to any traffic from the specified interface.
	To override the default profile for a specific tunnel, click the Tunneled Traffic tab, click Add , and then click individual entries to configure the following settings:
	• Tunnel Interface—Select the tunnel interface on the firewall.
	 QoS Profile—Select the QoS profile to apply to the specified tunnel interface.
	To remove a clear text or tunneled traffic entry, select the check box for the entry and click Remove .

Table 141. QoS Settings (Continued)

Defining QoS Profiles

Network > Network Profiles > QoS Profiles

For each interface, you can define QoS profiles that determine how the QoS traffic classes are treated. You can set overall limits on bandwidth regardless of class and also set limits for individual classes. You can also assign priorities to different classes. Priorities determine how traffic is treated in the presence of contention.



Note: Refer to "Configuring QoS for Firewall Interfaces" on page 356 for information on configuring firewall interfaces for QoS and refer to "Defining QoS Policies" on page 359 to configure the policies that will activate the QoS restrictions.

Field	Description
Profile Name	Enter a name to identify the profile (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Egress Max	Enter the maximum bandwidth allowed for this profile (Mbps).
Egress Guaranteed	Enter the bandwidth that is guaranteed for this profile (Mbps).
Classes	Specify how to treat individual QoS classes. You can select one or more classes to configure:
	• Class —If you do not configure a class, you can still include it in a QoS policy. In this case, the traffic is subject to overall QoS limits. Traffic that does not match a QoS policy will be assigned to class 4.
	• Priority —Click and select a priority to assign to this class. These are prioritized in the order listed (highest first):
	– real-time
	– high
	– medium
	- low
	• Egress Max—Click and enter a value (Mbps) for this class.
	• Egress Guaranteed—Click and enter a value (Mbps) for this class.
	When contention occurs, traffic that is assigned a lower priority is dropped. Real-time priority uses its own separate queue.

Table 142. QoS Profile Settings

Defining QoS Policies

Policies > QoS

The QoS policy determines how traffic is classified for treatment when it passes through an interface with QoS enabled. For each rule, you specify one of eight classes. You can also assign a schedule to specify which rule is active. Unclassified traffic is automatically assigned to class 4.



Note: Refer to "Configuring QoS for Firewall Interfaces" on page 356 for information on configuring firewall interfaces for QoS and refer to "Defining QoS Profiles" on page 358 for information on configuring classes of service.

To view just the rules for a specific virtual system, select the system from the **Virtual System** drop-down list and click **Go**. To apply a filter to the list, select from the **Filter Rules** drop-down list. To view just the rules for specific zones, select a zone from the **Source Zone** and/or **Destination Zone** drop-down lists, and click **Filter by Zone**.



Note: Shared polices pushed from Panorama are shown in green and cannot be edited at the device level.

To add a new QoS rule, do one of the following:

- Click **Add** at the bottom of the page and configure the rule. A new rule is added to the bottom of the list.
- Select **Clone Rule**, or select a rule by clicking the white space of the rule, and select **Clone** at the bottom of the page (a selected rule has a yellow background). The copied rule is inserted below the selected rule.

Field	Description
General Tab	
Name	Enter a name to identify the rule (up to 31 characters). The name is case- sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Description	Enter an optional description.
Tag	If you need to tag the policy, click Add to specify the tag.
	A policy tag is a keyword or phrase that allows you to sort or filter policies. This is useful when you have defined many policies and want to view those that are tagged with a particular keyword. For example, you may want to tag certain security policies with Inbound to DMZ, decryption policies with the words Decrypt and No-decrypt, or use the name of a specific data center for policies associated with that location.
Source Tab	

Field	Description
Source Zone	Select one or more source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.
Source Address	Specify a combination of source IPv4 or IPv6 addresses for which the identified application can be overridden. To select specific addresses, choose select from the drop-down list and do any of the following:
	• Select the check box next to the appropriate addresses and/or address groups in the Available column, and click Add to add your selections to the Selected column.
	• Enter the first few characters of a name in the Search field to list all addresses and address groups that start with those characters. Selecting an item in the list will set the check box in the Available column. Repeat this process as often as needed, and then click Add .
	• Enter one or more IP addresses (one per line), with or without a network mask. The general format is:
	 To remove addresses, select the appropriate check boxes in the Selected column and click Remove, or select any to clear all addresses and address groups.
	To add new addresses that can be used in this or other policies, click New Address (refer to "Defining Applications" on page 233). To define new address groups, refer to "Defining Address Groups" on page 230.
Source User	Specify the source users and groups to which the QoS policy will apply.
Negate	Select the check box to have the policy apply if the specified information on this tab does NOT match.
Destination Tab	
Destination Zone	Select one or more source zones (default is any). Zones must be of the same type (Layer 2, Layer 3, or virtual wire). To define new zones, refer to "Defining Security Zones" on page 151.
Destination Address	Specify a combination of source IPv4 or IPv6 addresses for which the identified application can be overridden. To select specific addresses, choose select from the drop-down list and do any of the following:
	 Select the check box next to the appropriate addresses and/or address groups and/or in the Available column, and click Add to add your selections to the Selected column.
	• Enter the first few characters of a name in the Search field to list all addresses and address groups that start with those characters. Selecting an item in the list will set the check box in the Available column. Repeat this process as often as needed, and then click Add .
	• Enter one or more IP addresses (one per line), with or without a network mask. The general format is:
	<ip_address>/<mask></mask></ip_address>
	• To remove addresses, select the appropriate check boxes in the Selected column and click Remove , or select any to clear all addresses and address groups.
	To add new addresses that can be used in this or other policies, click New Address (refer to "Defining Applications" on page 233). To define new address groups, refer to "Defining Address Groups" on page 230.

Table 143. QoS Rule Settings (Continued)
Field	Description
Negate	Select the check box to have the policy apply if the specified information on this tab does NOT match.
Application Tab	
Application	Select specific applications for the QoS rule. To define new applications, refer to "Defining Applications" on page 233. To define application groups, refer to "Defining Application Groups" on page 238.
	If an application has multiple functions, you can select the overall application or individual functions. If you select the overall application, all functions are included, and the application definition is automatically updated as future functions are added.
	If you are using application groups, filters, or container in the QoS rule, you can view details on these objects by holding your mouse over the object in the Application column, click the down arrow and select Value . This enables you to easily view application members directly from the policy without having to go to the Object tabs.
Service/ URL Category Tab	
Service	Select services to limit to specific TCP and/or UDP port numbers. Choose one of the following from the drop-down list:
	 any—The selected applications are allowed or denied on any protocol or port.
	 application-default—The selected applications are allowed or denied only on their default ports defined by Palo Alto Networks. This option is recommended for allow policies.
	• Select—Click Add. Choose an existing service or choose Service or Service Group to specify a new entry. Refer to "Services" on page 239 and "Service Groups" on page 240.
URL Category	Select URL categories for the QoS rule.
	 Choose any to ensure that a session can match this QoS rule regardless of the URL category.
	• To specify a category, click Add and select a specific category (including a custom category) from the drop-down list. You can add multiple categories. Refer to "Custom URL Categories" on page 242 for information on defining custom categories.
Other Settings Tab	
Class	Choose the QoS class to assign to the rule, and click OK . Class characteristics are defined in the QoS profile. Refer to "Defining QoS Profiles" on page 358 for information on configuring settings for QoS classes.
Schedule	Choose the calendar icon to set a schedule for the QoS policy to apply.

Table 143.	QoS Rule	Settings	(Continued)
------------	----------	----------	-------------

Displaying QoS Statistics

► Network > QoS

The table on the **QoS Policies** page indicates when QoS is enabled, and includes a link to display QoS statistics. An example is shown in the following figure.

Figure 46. QoS Statistics

QoS Statistics												0 🖂
Name	Guaranteed Egress (Mbps)	Maximum Egress (Mbps)	Runtime Bandwidth (Mbps)	Ban	dwi	dth Ap	plications	Source Users	Destination Users	Security Rules	QoS Rules	
ethernet1/2			0	~	clas	is 5	Class 6		s 7 🗹 class 8			
🖃 🚍 default-group	0	999	0									
E class 1	0	999	0		100	· · · · ·						
E class 2	0	999	0									
E class 3	0	999	0									
E class 4	0	999	0									
E class 5	0	999	0									
E class 6	0	999	0		80	1						
\Xi class 7	0	999	0									
E class 8	0	999	0									
🖃 😋 tunnel-traffic			0	P								
🕀 🦳 tunnel.200	0.001	1000	0	na i	60							
i 🔁 🧰 bypass-traffic	0	0	0	me Bandwidth (Mbps)	40							_
					20				Time			
Note: Bandwidth limits shown include	e hardware adjustm	ient factor.										

The left panel shows the QoS tree table, and the right panel shows data in the following tabs:

- **QoS Bandwidth**—Shows the real time bandwidth charts for the selected node and classes. The information is updated every two seconds.
- Session Browser—Lists the active sessions of the selected node and/or class.
- Application View—Lists all active applications for the selected QoS node and/or class.

Chapter 11 Setting Up a VM-Series Firewall

This chapter describes how to install the VM-Series firewall:

- "Overview" in the next section
- "System Requirements and Limitations" on page 364
- "About Licensing the VM-Series Firewall" on page 365
- "Installing and Licensing the VM-Series Firewall" on page 366
- "Troubleshooting" on page 370

Overview

The Palo Alto Networks VM-Series firewall is the virtualized form of the Palo Alto Networks firewall. It is positioned for use in a virtualized data center environment and is particularly well suited for private and public cloud deployments. By using virtual machine technology, you can install this solution on any x86 device that is capable of running VMware ESXi, without the need to deploy Palo Alto Networks hardware.

The VM-Series firewall has many characteristics in common with the Palo Alto Networks hardware firewalls, including common features and management interfaces. The main difference is in the deployment methods used in a virtual environment in order to fit your needs. Once the virtual firewall is installed, you will use it and manage it much like you do with the hardware firewalls.

The VM-Series firewall is distributed using the Open Virtualization Format (OVF), which is a standard method of packaging and deploying virtual machines.

Note: This section covers the basic installation steps needed to deploy a VM-Series firewall. For more details and uses cases, refer to the VM-Series firewall Tech Note at *https://live.paloaltonetworks.com/community/documentation.*

System Requirements and Limitations

This section lists requirements and limitations for the VM-Series firewall.

Requirements

You can create and deploy multiple instances of the VM-Series firewall on an ESX(i) server. Because each instance of the firewall requires a minimum resource allocation — number of CPUs, memory and disk space— on the ESX(i) server, make sure to conform to the specifications below to ensure optimal performance.

The VM-Series firewall has the following requirements:

- VMware ESX(i) with vSphere 4.1 and 5.0.
- Minimum of two vCPUs per VM-Series firewall. One will be used for the for the management plane and one for the dataplane. You can add up to eight additional vCPUs for the dataplane in the following increments: 2, 4, or 8 vCPUs.
- Minimum of two network interfaces (vmNICs). One will be a dedicated vmNIC for the management interface and one for the data port. You can then add up to eight more vmNICs for data traffic.
 The VM-Series firewall requires that promiscuous mode is set to "accept" on the port group of the virtual switch to which the data ports on the firewall are attached.
- Minimum of 4GB of memory. Any additional memory will be used by the management plane only.
- Minimum of 40GB of virtual disk space. You can add an additional disk of up to 2TB for logging purposes.

Limitations

The VM-Series firewall functionality is very similar to the Palo Alto Networks hardware firewalls, but with the following limitations:

- Dedicated CPU cores are required.
- Only High Availability (HA) lite is supported (active/passive with no stateful failover).
- High Availability (HA) "Link Monitoring" is only supported on VMware ESXi installations that support DirectPath I/O.
- Up to 10 total ports can be configured; this is a VMware limitation. One port will be used for management traffic and up to 9 can be used for data traffic.
- Only the vmxnet3 driver is supported.
- Virtual systems are not supported.
- vMotion is not supported.
- Jumbo frames are not supported.
- Link Aggregation is not supported.

About Licensing the VM-Series Firewall

The VM-Series firewall is available in three models — VM-100, VM-200, and VM-300— and each model is licensed for a maximum capacity. Capacity is defined in terms of the number of sessions, rules, security zones, address objects, IPSec VPN tunnels and SSL VPN tunnels that the VM-Series firewall is optimized to handle. When purchasing a license, make sure to purchase the correct model for your network requirements. While each model has a distinct orderable part number, the software package that is used to create the VM-Series firewall is the same across all models.

When you purchase a VM-Series firewall, you will receive a set of auth-codes over email. Typically the email includes a capacity auth-code for the model purchased (VM-100, VM-200, VM300), a software and support auth-code (for example, PAN-SVC-PREM-VM-100 SKU auth-code) that provides access to software/content updates and support. If you purchased additional subscriptions for Threat Prevention, URL Filtering, GlobalProtect, or WildFire, a list of the other auth-codes purchased with the order are included.

If you do not have an existing support account, you must use the capacity auth-code to register and create an account on the Support Portal. After your account is verified and the registration is complete, you will be able to log in and download the software package required to install the VM-Series firewall. If you have an existing support account, you can access the "VM-Series Authentication Code" link on the support software page to manage your VM-Series firewall licenses and download the software.



Note: If you have an evaluation copy of the VM-Series firewall and would like to convert it to a fully licensed (purchased) copy, clone your VM-Series firewall and use the instructions in the following sections to register and license the purchased copy of your VM-Series firewall.

Registering the VM-Series Firewall

You are required to register your capacity auth-code with your support account.

- 1. Log in to https://support.paloaltonetworks.com with your account credentials.
- 2. Select **My VM-Series Auth-Codes** link on the home page and click **Add VM-Series Auth-Code**.
- 3. Enter the capacity auth-code you received by email, and click the **Add VM-Series Auth-Code** button. The page will refresh and you will see the list of auth-codes registered to your support account.
- 4. Continue with "Installing and Licensing the VM-Series Firewall" on page 366.

Installing and Licensing the VM-Series Firewall

This section describes the steps needed to install and license the VM-Series firewall.

Installing a VM-Series firewall:

1. Obtain a VM-Series license then download the Open Virtualization Format (OVF) template from *https://support.paloaltonetworks.com* software site.

The OVF is downloaded as a zip archive that is expanded into three files:

- OVF extension—The OVF descriptor file that contains all metadata about the package and its contents.
- MF extension—The OVF manifest file that contains the SHA-1 digests of individual files in the package.
- VMDK extension—The disk image file that contains the virtualized version of the PAN-OS firewall.



Note: The above files contain the base installation. After the base installation is complete, you will need to download and install the latest version of the VM-Series firewall software from the support site. This will ensure that you have the latest fixes that were implemented since the base OVF file was created.

2. Before deploying the template, it is helpful to setup virtual standard switch(es) and virtual distributed switch(es) that you will need for the VM-Series firewall. The VM-Series firewall requires that any attached virtual switch has promiscuous mode enabled.

To configure a virtual standard switch:

- a. Configure a virtual standard switch from the vSphere Client by navigating to Home > Inventory > Hosts and Clusters.
- b. Click the **Configuration** tab and under **Hardware** click **Networking**. For each VM-Series firewall attached virtual switch, click on **Properties**.
- c. Highlight the virtual switch and click **Edit**. In the vSwitch properties, click the **Security** tab and set **Promiscuous Mode:** to **Accept** and then click **OK**. This change will propagate to all port groups on the virtual switch.

To configure a virtual distributed switch:

- a. Navigate to **Home > Inventory > Networking**. Highlight the **Distributed Port Group** you want to edit and select the **Summary** tab.
- b. Click on Edit Settings and select Policies > Security and set Promiscuous Mode: to Accept and then click OK.
- 3. Deploy the OVF template:
 - a. Log in to vCenter using the vSphere client. You can also go directly to the target ESXi host if needed.
 - b. From the vSphere client, select **File > Deploy OVF Template**.

- c. Browse to the OVF template that you downloaded in step 1, select the file and then click **Next**. Review the templates details window and then click **Next** again.
- d. Name the VM-Series firewall instance and in the **Inventory Location:** window, select a Data Center and folder and click **Next**.

Source	Name:
Name and Location	Phoenix
 Host / Cluster Resource Pool Disk Format Ready to Complete 	The name can contain up to 80 characters and it must be unique within the inventory folder. Inventory Location:
	WIN-A19VCGVNGV2 WiN-A19VCGVNGV2 Mit-virt-lab per host use case Sniffers Templates QA
Help	< Back Next > Cancel

- e. Select an ESXi host for the VM-Series firewall and click Next.
- f. Select the datastore to use for the VM-Series firewall and click Next.
- g. Leave the default settings for the datastore provisioning and click **Next**. The default is "Thick Provision Lazy Zeroed".
- h. Select the networks to use for the two initial vmNICs. The first vmNIC will be used for the management interface and the second vmNIC for the first data port. Make sure that the **Source Networks** is mapped to the correct **Destination Networks**.

Source OVF Template Details	Map the networks used in this OVF	template to networks in your inventory	
Host / Cluster	Source Networks	Destination Networks	_
Storage	VM Network	VM Network	
Disk Format	VM Network 2	VM Network	*
Network Mapping		VM Network	
Neday to complete		dvPG301	
		dvPG3021/C	
	Description:		
	The VM Network 2 network		*
			~
	Warning: Multiple source networks	are mapped to the host network: VM Network	
Help		< Back Next >	Cancel

i. Review the details window, click the **Power on after deployment** check box and then click **Next**.

You can monitor the **Recent Tasks** list to view the progress of the deployment. When the deployment is complete, click the **Summary** tab to review the current status.

- 4. To perform the initial configuration of the VM-Series firewall, follow these steps:
 - a. The default management IP is 192.168.1.1. To change this, launch the vSphere console, click the **Summary** tab and under **Commands** click **Open Console**. You can also right click the VM and select Open Console.
 - b. Log in with the default login user name **admin** password **admin** and then type **configure** to enter configuration mode.
 - c. Set the desired IP, netmask, gateway, and DNS IP. Example: set deviceconfig system ip-address 10.1.1.5 netmask 255.255.255.0 default-gateway 10.1.1.1 dns-setting servers primary 10.0.0.245
 - d. Type **commit** to make the changes active.
 - e. Test network connectivity to your default gateway and a known server. Example: ping host 10.1.1.1 or ping host 10.0.0.245.
- 5. Activate the license. Pick one of the following options based on whether or not the VM-Series firewall has direct internet access.
 - a. If your firewall has direct internet access:
 - i. Navigate to the **Device >Licenses** tab and select the **Activate feature using authentication code** link.
 - ii. Enter the capacity auth-code that you registered on the Support website. The firewall will connect to the update server (updates.paloaltonetworks.com), and download the license and reboot automatically.
 - iii. Log back in to the web interface and confirm that the **Dashboard** displays a valid serial number. If the term *Unknown* displays, it means the device is not licensed.
 - iv. On the **Device->** Licenses tab, verify that "PA-VM" license is added to the device.

- b. In case the ESXi server that your VM-Series firewall is installed on does not have internet access:
 - i. Navigate to the **Device > Licenses** tab and click the **Activate Feature using Auth Code** link.
 - ii. Click Download Authorization File, and download the *authorizationfile.txt* on the client machine.
 - iii. Copy the *authorizationfile.txt* to a computer that has access to the internet and log in to the support portal. Click My VM-Series Auth-Codes link and select the applicable auth-code from the list and click the Register VM link.
 - iv. On the **Register VM Device** tab upload the authorization file. This will complete the registration process and the serial number of your VM-Series firewall will be attached to your account records.
 - v. Navigate to the **My Devices** tab and search for the VM-Series device just registered and click the **PA-VM** link. This will download the VM-Series license key to the client machine.
 - vi. Copy the license key to the machine that can access the web interface of the VM-Series firewall and navigate to the **Device -> Licenses** tab.
 - vii.Click **Manually Upload License** link and enter the license key. The license will be activated on the device and the device will reboot.
 - viii.Log in to the device and confirm that the **Dashboard** displays a valid serial number and that the PA-VM license displays in the **Device > Licenses** tab.



Note: In order to complete the license installation process, you must reboot your VM-Series firewall instance.

6. Now that the VM-Series firewall has network connectivity and the base PAN-OS software is installed, you need to upgrade to the latest version of PAN-OS (a support license is required).



Note: The base VM-Series firewall software is installed with the OVF file. It's recommended that you upgrade to the latest content update and PAN-OS release to ensure you have all of the latest fixes. At each feature release, a new base image will be created as part of the OVF file.

- a. From the web interface, navigate to **Device > Licenses** and make sure you have the correct VM-Series firewall license and that the license is activated.
- b. To upgrade the VM-Series firewall PAN-OS software, navigate to **Device > Software**.
- c. Click **Refresh** to view the latest software release and also review the **Release Notes** to view a description of the changes in a release and to view the migration path to install the software.
- d. Click Download to retrieve the software then click Install to install it.

For details on downloading and installing PAN-OS software, refer to "Upgrading/ Downgrading the PAN-OS Software" on page 50.

Troubleshooting

Many of the troubleshooting steps for the VM-Series firewall are very similar to the hardware versions of PAN-OS. When problems occur, you should check interface counters, system log files, and if necessary, use debug to create captures. For more details on PAN-OS troubleshooting, refer to the Packet Based Troubleshooting Tech Note at *https://live.paloaltonetworks.com/community/documentation.*

Similar to a physical environment, it is sometimes useful to have a separate troubleshooting client to capture traffic in the virtualized environment. It can be helpful to build a fresh OS from scratch with common troubleshooting tools installed, such as: tcpdump, nmap, hping, traceroute, iperf, tcpedit, netcat, etc. Going forward, each time it is needed, the troubleshooting client can be quickly deployed to the virtual switch(es) in question and used to isolate networking problems.

For performance related issues on the firewall, first check the **Dashboard** from the firewall web interface. For information in the VM server, in vSphere Client go to **Home > Inventory > VMs and Templates**, select the VM-Series firewall instance and click the **Summary** tab. Under **Resources**, check the statistics for consumed memory, CPU and storage. For resource history, click the **Performance** tab and monitor resource consumption over time.

You can also view alerts or create a tech support or stats dump files from **Device > Support**. For more information, refer to "Viewing Support Information" on page 117.

Chapter 12 Setting Up Panorama

This chapter describes how to set up the Panorama centralized management system:

- "Overview" in the next section
- "Setting Up Panorama as a Virtual Appliance" on page 372
- "Setting up Panorama on an M-Series Appliance" on page 376
- "Configuring High Availability (HA)" on page 377



Note: For information on using Panorama, see "Central Device Management Using Panorama" on page 381.

Overview

Panorama is the centralized management system for the Palo Alto Networks family of nextgeneration firewalls, and is available both as a dedicated hardware platform and as a VMware virtual appliance that meets your server consolidation needs.

Panorama provides centralized visibility and management for all the devices on your network. Because Panorama shares the same web-based look and feel as the individual device interface, you can seamlessly transition in to managing the devices centrally and reducing the administrative effort in managing multiple devices.

For information on installing Panorama on VMware ESX(i) 3.5 or later, see "Setting Up Panorama as a Virtual Appliance" on page 372

For information on setting up the hardware based M-100 management appliance, see "Setting up Panorama on an M-Series Appliance" on page 376.

Setting Up Panorama as a Virtual Appliance

The pre-requisites for setting up a Panorama virtual appliance are:

- System requirements
 - VMware ESX(i) 3.5 or later
 - Quad Core CPU (3GHz) (use 4 GHz if you have 10 or more active firewalls)
 - 2-4 GB RAM (use 4 GB if you have 10 or more active firewalls)
 - 34 GB disk space
 - VMware vSphere Client 4.x or VMware Infrastructure Client 3.5
- Use the assigned serial number to register Panorama on the support site at *https://support.paloaltonetworks.com*. After you register Panorama on the support site, you will have access to the Panorama software downloads page.
- Download the latest Panorama base image zip file to the server on which you will be installing Panorama. The virtual appliance installation procedure uses the Open Virtual Machine Format (OVF) template file, which is included in the base image.

Installing Panorama

Follow these steps to install Panorama on your ESX(i) server:

- 1. Unzip the Panorama zip file to find the *panorama-esx.ovf* template file for installation.
- 2. Open the VMware vSphere Client and connect to your VMware server from the login screen.
- 3. Choose **File > Deploy OVF Template**.
- 4. Browse to select the *panorama-esx.ovf* file from the recently unzipped Panorama base image, and click **Next**.
- 5. Confirm that the product name and description match the downloaded version, and click **Next**.
- 6. Choose a name for the Panorama image, and click Next.
- 7. Select a datastore location to install the Panorama image, and click Next.
- 8. If prompted, choose **Thick provisioned format** for the disk format, and click **Next**.
- 9. Confirm the options you selected and then click Finish to begin the installation process.
- 10. When the installation is complete, choose the newly installed Panorama image and click the **Power On** button.

When the Panorama virtual machine boots, the installation process is complete. Continue with the next section to use the console and perform initial set up.

Configuring the Panorama Network Interface

To configure the Panorama network interface using the Panorama virtual machine console on your ESX(i) server:

- 1. Log in to the CLI. Enter **admin** for both the username and password fields.
- 2. At the prompt, type configure to switch to the configuration mode.
- 3. Define the network access configuration for the management interface. Enter the following commands on one line. set deviceconfig system ip-address <Panorama-IP> netmask <netmask> default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>

where *<Panorama-IP>* is the IP address, *<netmask>* is the subnet mask, *<gateway-IP>* is the IP address of the network gateway, and *<DNS-IP>* is the IP address of the Domain Name System (DNS) server.

- 4. Type **commit** to save the changes and activate them.
- 5. Type **Exit** to leave the configuration mode.
- 6. Test network connectivity to your default gateway or another server (*<target-IP>*). ping host <target-IP>

Make sure that you can successfully ping the gateway and the Internet.

Next Steps

- To log in to Panorama, and to change the default password, see "Logging in to Panorama" on page 377.
- To set up additional log storage capacity for your Panorama virtual appliance, see "Expanding the Log Storage Capacity" on page 373.
- To configure high availability, see "Configuring High Availability (HA)" on page 377.
- To start managing devices using Panorama, see the following sections:
 - To add devices, see "Panorama Administrator Roles, Profiles, and Accounts" on page 387.
 - Top verify that each managed device is configured with the IP address of the Panorama server, see "System Setup, Configuration, and License Management" on page 30.

Expanding the Log Storage Capacity

By default, Panorama maintains internal storage for log files and statistical data. The default Panorama installation is set up with a single disk partition for all data; 10 GB of space is allocated for log storage on the partition. To support environments that require more storage space, you can select from two options:

• Create a custom virtual disk up to 2TB for ESX or ESXi. For instructions, see "Adding a Virtual Disk" on page 374.

or,

• Configure an external NFS data store. For instructions, see "Setting Up Storage Partitions" on page 374.

Adding a Virtual Disk

If you require more storage space than the 10GB that is provided by default on the Panorama device, use the following steps to create a custom virtual disk:

- 1. On your ESX(i) server, select the Panorama virtual machine.
- 2. Click Edit Settings.
- 3. Click Add to launch the Add Hardware wizard.
- 4. Choose Hard Disk from the list of hardware types and click Next.
- 5. Choose the Create a new virtual disk option and click Next.
- 6. Choose SCSI for the Virtual Disk Type and click Next.
- 7. Select **Specify a datastore** in the location field and enter a name and path or select using the **Browse** button.
- 8. Click Finish.

The new disk is shown in the list of devices for the virtual machine.

9. Start the Panorama virtual machine.

On the first start after adding the new disk, Panorama initializes the new disk for use. This process may take several minutes to a few hours, depending on the size of the newly added disk.

After the system starts with the new disk, any existing logs on the default disk are moved to the new virtual disk, and all future log entries are written to the new disk. If the virtual disk is removed, Panorama automatically reverts back to logging to the default internal 10 GB disk.

If you have already added a virtual disk and would like to replace it with a larger or different virtual disk, you must first remove the installed virtual disk. However, when a first virtual disk is removed you can no longer access the logs on that disk.



Note: To allow for redundancy, use the virtual disk in a RAID configuration. RAID 10 provides the best write performance for applications with high logging characteristics. For further performance improvements, optimize the drives for sequential writing of a small number of large files.

Setting Up Storage Partitions

Panorama > Setup > Storage Partition Setup

To configure an external NFS data store:

Click the **Storage Partition Setup** link on the Panorama **Setup** page, and specify the following settings.



Note: You must reboot the Panorama server after configuring the storage partition settings.

Table 144. Storage Partition Settings

Field	Description
Internal	Maintains storage space for log files and statistical date on the Panorama device.
NFS v3	Specifies an external NFS server mount point for storage. Configure the following settings:
	• Server —Specify the fully qualified domain name (FQDN) or IP address of the NFS server.
	• Log Directory—Specify the full path name of the directory where the logs will be stored.
	• Protocol —Specify the protocol for communication with the NFS server (UDP or TCP).
	• Port—Specify the port for communication with the NFS server.
	• Read Size —Specify the maximum size (bytes) for NFS read operations (range 256 - 32768).
	• Write Size—Specify the maximum size (bytes) for NFS write operations (range 256 - 32768).
	• Copy On Setup —Select the check box to mount the NFS partition and copy any existing logs to the destination directory on the server when the Panorama device boots.
	 Test Logging Partition—Click to perform a test that mounts the NFS partition and presents a success or failure message.

Setting up Panorama on an M-Series Appliance

The hardware-based Panorama solution uses the M-100 device. Refer to the *M-100 Hardware Reference Guide* for information on rack-mounting the appliance and powering it on.

Complete the following tasks before you begin with performing initial setup.

- Register your M-100 device at http://support.paloaltonetworks.com to obtain the latest software updates and to activate support.
- Obtain an IP address for Panorama from your system administrator.
- Set your computer's IP address to 192.168.1.2 and the subnet mask to 255.255.255.0.

Performing Initial Setup

The initial setup allows you to use the default IP address to access the Panorama Management Console and then assign an IP address on the management interface so that you can manage the device across your network.

To perform initial setup:

- 1. Connect your computer to the management port (MGT) using an RJ-45 Ethernet cable (provided).
- 2. Power on your computer.
- 3. Log in to Panorama.
 - a. Launch an Internet browser on your computer and enter https://192.168.1.1.
 - b. Type admin in both the **Name** and **Password** fields.
 - c. Click Login.
- 4. Click **Panorama > Setup**. Click the small gear icon (Edit...) on the **Management Interface Settings** table.
- 5. Enter the new IP address and related network access information for using the management interface (MGT) on your enterprise management network. Click **OK**.
- 6. Commit your changes to the device. Click **Commit** and select Panorama as the **Commit Type**, then click **OK**.
- 7. Disconnect your computer from the M-100 device.
- 8. Connect the MGT port on the front panel of the M-100 to the enterprise management network. Your device is now set up for access across your network.

Next steps:

- To log in and verify access to the Panorama Management Console. See "Logging in to Panorama" on page 377.
- To change the default login credentials, see "Changing the Default Password" on page 377.

Logging in to Panorama

To log in to the web interface of Panorama:

- 1. Launch a web browser and enter **https://<Panorama IP address>**. The browser automatically opens the Palo Alto Networks login page.
- 2. Enter **admin** in both the **Name** and **Password** fields.
- 3. Click Login.

Changing the Default Password

Panorama is configured with a default password. To improve security, you are encouraged to strengthen the password on first-time login.

- 1. Choose **Panorama > Administrators > admin**.
- 2. Enter admin in the Old Password field.
- 3. Enter a new password (case-sensitive, up to 15 characters) in the **New Password** field and re-enter the password in the **Confirm New Password** field.
- 4. Click OK.
- 5. Click **Commit** and select **Panorama** as the **Type**.
- 6. Click OK.

Configuring High Availability (HA)

Panorama > High Availability



Note: HA is supported only for managed devices running Release 4.0 or later. It is not backward compatible with Release 3.1 or earlier.

High availability (HA) allows for redundancy in the event of a device failure. For ensuring HA, you can deploy a pair of hardware-based Panorama appliances or a pair of Panorama virtual appliances in a HA peer configuration that provide synchronized connections to the managed firewalls. Among the peers on the HA configuration, one device must be designated as active and the other as passive. The peers maintain a heartbeat, or a periodic ICMP ping, to verify operational status. If the active Panorama device becomes unavailable, the passive server takes over temporarily. With preemption enabled, the default setting, when the active device becomes available again, the passive device relinquishes control and returns to the passive state.



Note: To configure a HA pair of Panorama virtual appliances, you must have two Panorama licenses with unique serial numbers for each virtual instance.

When configuring HA for Panorama virtual appliances, you also need to designate a priority level of either primary or secondary for each peer in the pair. This primary or secondary configuration determines which peer is designated as the primary recipient for logs sent by the managed firewalls. You can configure Panorama to use the same log external storage facility for the assigned primary and secondary devices (Network File System or NFS option) or configure logging internally. If you use the NFS option, only the primary recipient receives the logs that are sent from the managed firewalls. However, if local logging is enabled, by default the logs are sent to both the primary and the secondary recipient.

To enable HA on Panorama, configure the followings settings.

Field	Description
Setup	
Enable HA	Select the check box to enable HA.
Peer HA IP Address	Enter the IP address of the MGT interface of the Panorama peer.
Enable Encryption	Enable encryption after exporting the HA key from the HA peer and importing it onto this device. The HA key on this device must also be exported from this device and imported on the HA peer. When enabled, the MGT interface encrypts communication between the HA peers. The key import/export is done on the Panorama > Certificate Management > Certificates page. See "Importing, Exporting and Generating Security Certificates" on page 86.
	<i>Note:</i> HA connectivity uses TCP port 28 with encryption enabled and 28769 when encryption is not enabled.
Monitor Hold Time (ms)	Enter the length of time (ms) that the system will wait before acting on a control link failure (1000-60000 ms, default 3000 ms).
Election Settings	
Priority	Choose Primary or Secondary .
Preemptive	Select the check box to enable the primary Panorama device to resume active operation after recovering from a failure. If this setting is off, then the secondary device remains active even after the higher priority device recovers from a failure.
Preemption Hold Time (min)	Enter the time a passive device will wait before taking over as the active device (range 1-60 min, default 1).
Promotion Hold Time (ms)	Enter the time that the secondary device will wait before taking over (range 0-60000 ms, default 2000).
Hello Interval (ms)	Enter the number of milliseconds between the hello packets sent to verify that the other device is operational (ranges 8000-60000 ms, default 8000).
Heartbeat Interval (ms)	Specify how frequently Panorama sends ICMP pings to the HA peer (range 1000-60000 ms, default 1000).
Monitor Fail Hold Up Time (ms)	Specify the interval that Panorama waits following a path monitor failure before attempting to re-enter the passive state (default 0 ms). During this period, the device is not available to take over for the active device in the event of failure.
Additional Master Hold Up Time (ms)	Specify the interval during which the preempting device remains in the passive state before taking over as the active device (default 7000 ms).

Table 145. Panorama HA Settings

Field	Description
Path Monitoring	
Enabled	Select the check box to enable path monitoring. Path monitoring enables Panorama to monitor specified destination IP addresses by sending ICMP ping messages to make sure that they are responsive.
Failure Condition	Select whether a failover occurs when any or all of the monitored path groups fail to respond.
Path Groups	Define one or more path groups to monitor specific destination addresses. To add a path group, specify the following and click Add :
	 Name—Specify a name for the path group.
	• Enabled—Select the check box to enable the path group.
	 Failure Condition—Select whether a failure occurs when any or all of the specified destination addresses fails to respond.
	• Ping interval —Specify a length of time between ICMP echo messages to verify that the path is up (range 1000-60000 ms, default 5000).
	• Destination IPs —Enter one or more destination addresses to be monitored (multiple addresses must be separated by commas).
	• Ping Interval —Specify the interval between pings that are sent to the desti- nation address (range 1000-60000 milliseconds, default 5000 milliseconds).
	• Ping Count —Specify the number of failed pings before declaring a failure (range 3-10 pings, default 3 pings).
	To delete a path group, select the group, and click Delete .

Table 145. Panorama HA Settings (Continued)

Switching the Logging Priority in an HA Pair

When a failover occurs for Panorama virtual appliances in an HA configuration with NFSbased logging, the logging functionality is interrupted. To re-enable logging following a failure, you must promote the secondary Panorama device to function as the primary so that it can establish a connection with the NFS-based log partition.



Note: For configurations that use internal logging instead of NFS, follow the instructions through Step 2 of the procedure in this section to switch the logging priority of the secondary device to primary.

For the following procedure, assume that the active primary is running on server S1 and the passive secondary is running on S2. Failover has occurred, and S2 has become the active secondary.

To assign S2 as the primary log recipient, follow these steps:

- 1. Power S1 off.
- 2. Configure S2 to be primary and commit the configuration:
 - a. Choose **Panorama > High Availability.**
 - b. Edit the election settings and change Priority from Secondary to Primary.
 - c. Commit the changes, rebooting the device when prompted. The reboot is required because the configuration refers to NFS storage.

3. Execute the CLI command request high-availability convert-to-primary.

If S1 is connected as the HA peer to S2 and NFS storage is specified, then the **convert-toprimary** command will fail, indicating that the HA peer (S1) needs to be powered down before the operation can succeed. If the peer is not connected, the system dynamically mounts the NFS disk, converts the ownership of the partition to S2, and unmounts the partition.

4. Reboot S2.

When S2 comes up, it is able to write to the NFS-based log partition.

Chapter 13 Central Device Management Using Panorama

This chapter describes how to use the Panorama centralized management system to manage multiple firewalls:

- "Accessing the Panorama Web Interface" in the next section
- "Using the Panorama Interface" on page 382
- "Adding Devices" on page 385
- "Panorama Administrator Roles, Profiles, and Accounts" on page 387
- "Specifying Panorama Access Domains for Administrators" on page 392
- "Working with Objects" on page 395
- "Working with Policies" on page 393
- "Templates" on page 399
- "Logging" on page 403
- "Viewing Firewall Deployment Information" on page 416
- "Backing Up Firewall Configurations" on page 417
- "Scheduling Configuration Exports" on page 417
- "Upgrading the Panorama Software" on page 419

Accessing the Panorama Web Interface

To access the Panorama interface for centralized firewall management, log in to the Panorama server web interface:

1. Launch your preferred web browser and enter https://Panorama IP address

The browser automatically opens the Palo Alto Networks login page.

2. Enter the login name and password and click **Login**.

Using the Panorama Interface

Panorama allows you to view information about multiple devices in your network and to manage devices from a central web interface.

To display information regarding the Palo Alto Networks firewalls in the network, the devices must be connected to the Panorama server.

Perform these steps to allow the devices to connect:

- 1. Add the IP address of the Panorama server to each device. Refer to "Defining Management Settings" on page 30.
- 2. Use the Panorama interface to add the devices. Refer to "Panorama Administrator Roles, Profiles, and Accounts" on page 387.

Certain Panorama configuration tabs will not appear until their respective components are configured. For instance, the Policies and Objects tab will only appear after adding device groups **Panorama > Device Groups** and **Device** and **Network** will only appear after adding templates from P**anorama > Templates**.

The Panorama tabs are described in the following table.

Page	Description
Dashboard	Displays general information about the managed devices, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available charts are displayed by default, but each user can remove and add individual charts, as needed.
ACC	Displays the overall risk and threat levels for the managed devices. Refer to "Using the Application Command Center" on page 253 and "Identifying Unknown Applications and Taking Action" on page 276.
Monitor	Allows you to view logs and reports. Refer to "Viewing Reports" on page 273.
Policies	Allows you to define policies to share across managed firewalls. Refer to "Logging and Reporting" on page 403 for information using the pages in this tab.
Objects	Allows you to define policy objects that are shared across the managed firewalls. Refer to "Logging and Reporting" on page 403.
Network	Allows you to apply network configuration options across managed firewalls using templates. Refer to "Templates" on page 399.

Table 146. Summary of Panorama Tabs

Page	Description
Device	Allows you to apply device configuration options across managed firewalls using templates. Refer to "Templates" on page 399.
Panorama	Allows you to configure Panorama and manage deployed firewalls. Refer to "Panorama Tab" in the next section.

Table 146. Summary of Panorama Tabs (Continued)

Panorama Tab

The **Panorama** tab is similar to the **Devices** tab for the firewall, but the settings apply to the Panorama device, not the managed firewalls. The following table describes the pages on this tab. To access a page, click the page name link on the side menu.

Table 147. Summary of Panorama Pages

Page	Description
Setup	Allows you to specify the Panorama host name, the network settings of the management interface, and the addresses of network servers (DNS and NTP). Refer to "Defining Management Settings" on page 30.
Templates	Allows you to create Templates that can be used to manage configuration options based on the Device and Network tabs, enabling you to deploy templates to multiple devices that have similar configurations. Refer to "Templates" on page 399.
Config Audit	Allows you to view and compare configuration files. Refer to "Defining Operations Settings" on page 37.
Managed Devices	Allows you to add devices for management by Panorama, push shared configuration to managed devices, and run comprehensive configuration audits on devices or entire device groups. Refer to "Adding Devices" on page 385.
Device Groups	Allows you to define sets of devices that are treated as a unit when creating objects and applying policies in Panorama. Refer to "Defining Device Groups" on page 386.
Managed Collectors	Allows you to configure and manage the log collector devices that will be used in your environment to distribute logging information for firewalls that are managed by Panorama. You can also use this pages to upgrade the software on your log collectors. You first download the latest Panorama software and you can then push the updated version to your log collectors by clicking Install on the Managed Collectors page.
	Note: Log collectors are comprised of the log collector software (part of the Panorama software package) and the M-100 hardware platform. The M-100 can be configured as a Panorama manager, a log collector, or both. The operational command to change the mode of an M-100 is request system logger-mode [panorama logger]. To view the current mode, run show system info match logger_mode. When an M-100 is in log collector mode, only the CLI is available for management.
	Refer to "Managing Log Collectors" on page 407.

Page	Description
Collector Groups	Allows you to group log collectors together so you can apply the same configuration settings to all collectors in the group. You also use the collector group to assign firewalls to log collectors.
	<i>Note:</i> You can add up to 4 log collector devices per collector group. Refer to "Defining Log Collector Groups" on page 411.
Admin Roles	Allows you to specify the privileges and responsibilities that are assigned to users who require access to Panorama. Refer to "Defining Administrator Roles" on page 58.
Password Profiles	Allows you to define password profiles, which can then be applied to Panorama administrators. You can configure the following profile options:
	 Required password change period (days)
	• Expiration warning period (days)
	Post Expiration Admin Login Count
	Post Expiration Grace Period (days)
Administrators	Allows you to define the accounts for users who require access to Panorama. Refer to "Creating Administrative Accounts" on page 59.
	Note: On the Administrators page for "," a lock icon is shown in the right column if an account is locked out. The administrator can click the icon to unlock the account.
High Availability	Allows you to configure a pair of Panorama devices to support high availability (HA). Refer to "Configuring High Availability (HA)" on page 377.
Certificate Management	Allows you to configure and manage certificate profiles, trusted certificate authorities, and OCSP responders. Refer to "Importing, Exporting and Generating Security Certificates" on page 86.
Log Settings	Allows you to define Simple Network Management Protocol (SNMP) trap sinks, syslog servers, and email addresses for distributing log messages. Refer to "Logging Configuration" on page 70.
Server Profiles	Allows you to specify profiles for servers that provide services to Panorama.
	Refer to the following sections:
	 "Configuring Email Notification Settings" on page 83
	 "Configuring SNMP Trap Destinations" on page 75
	 "Configuring Syslog Servers" on page 76
	 "Configuring RADIUS Server Settings" on page 65
	 "Configuring LDAP Server Settings" on page 66
	 "Configuring Kerberos Settings (Native Active Directory Authentication)" on page 67.
	• "Configuring Netflow Settings" on page 85
Authentication Profile	Allows you to specify a profile to authentication access to Panorama. Refer to "Authentication Profiles" on page 62.
Authentication Sequence	Allows you to specify sets of authentication profiles to use for access to Panorama. Refer to "Authentication Sequence" on page 67.
Access Domain	Provides the capability to limit administrators access to Device Groups, Templates, and the device contexts that administrators can switch to based on profiles. Refer to "Panorama Administrator Roles, Profiles, and Accounts" on page 387.

Table 147. Summary of Panorama Pages (Continued)

Page	Description
Scheduled Config Export	Allows you to collect running configurations from managed devices and deliver them daily to a File Transfer Protocol (FTP) server or by using Secure Copy (SCP) to securely transfer data between the Panorama server and a remote host. Refer to "Scheduling Configuration Exports" on page 417.
Software	Allows you to view the available Panorama software releases and download and install a selected software version. Refer to "Upgrading the Panorama Software" on page 419.
Dynamic Updates	Allows you to view the latest application definitions and information on new security threats, such as antivirus signatures (threat prevention license required) and update Panorama with the new definitions. Refer to "Updating Threat and Application Definitions" on page 55.
Support	Allows you to access product and security alerts from Palo Alto Networks. Refer to "Viewing Support Information" on page 117.
Deployment	Allows you to view current license information on the managed devices and install software, clients, and dynamic content on the devices. Refer to "Viewing Firewall Deployment Information" on page 416.
Master Key and Diagnostics	Allows you to specify a master key to encrypt private keys on the firewall. Private keys are stored in encrypted form by default even if a new master key is not specified. Refer to "Encrypting Private Keys and Passwords on the Firewall" on page 90.

Table 147. Summary of Panorama Pages (Continued)

Adding Devices

Panorama > Managed Devices

The **Managed Devices** page allows you to create a list of devices for centralized management. If devices are part of an HA pair, you must add both devices or virtual systems of the peers (if in multi-virtual system mode) to the same device group, and Panorama must push the configuration to both HA peer devices at the same time. If you target a rule to specific firewalls that are in an HA configuration, make sure to include both firewalls in the target selection.



Note: Panorama can manage PAN-OS devices running the same major release or earlier supported versions, but not devices running a later release version. For example, Panorama 4.0 can manage PAN-OS devices running 4.0 or earlier supported versions, but it cannot manage PAN-OS devices running 4.1.



Note: Managed devices communicate with Panorama using SSL through TCP port 3978.

To add devices:

- 1. Under the **Panorama** tab, click **Managed Devices** to open the **Managed Devices** page.
- 2. Click Add to open an editing window.
- 3. Enter the serial number of the device to be added, and click Add.

- 4. Add additional devices, as needed.
- 5. Click **OK**. The window closes and the **Managed Devices** page refreshes to show the newly added devices.
- 6. To delete a device:
 - a. Select the check box next to the device in the Managed Devices table.
 - b. Click Delete.
 - c. Click OK.
- 7. Select the **Group HA Peers** check box to group devices in high availability (HA) mode together.

This option allows you to easily identify devices that are in HA mode. When pushing shared policies, you can push to the grouped pair, instead of each device individually. Also, when adding a new device in Managed Devices, if they are in HA mode, both devices will be displayed together, so you can add both devices.

When viewing an HA pair, if the configuration does not match, a warning indicator will appear. You will also see an indicator if the HA devices are in different device groups.

This option is also independent for each section, so enabling and disabling in one area, will not enable/disable for all areas. The Group HA Peers option is present in the following Panorama areas:

- Managed Devices
- Templates
- Device Groups
- Policies tab (Target tab for all policy types)
- Commit dialog

Defining Device Groups

Panorama > Device Groups

Device groups are used to manage shared policies and objects. You can define device groups that consist of firewalls and/or virtual systems that you want to manage as a group, such as the firewalls that manage a group of branch offices or individual departments in a company. Each group is treated as a single unit when applying policies in Panorama.

You can add each device to at most one device group. Because virtual systems are considered distinct entities in Panorama, you can assign virtual systems within a device to different device groups.

The **Device Groups** page lists the device groups along with the information listed in the following table.

Field	Description			
Device Group Name	Enter a name to identify the group (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.			
Description	Enter a description for the group.			
Devices	Select the check box next to the device in the Managed Devices table and click Move , choose the new group and click OK .			
Master Device	Select a device to use as the master. The master device is the firewall from which Panorama gathers User-ID information for use in policies. The gathered user and group mapping information is specific to a device group and can come from only one device (the master) inside the group.			
Group HA Peers	Select the check box to group devices in high availability (HA) mode together.			
	This option allows you to easily identify devices that are in HA mode. When pushing shared policies, you can push to the grouped pair, instead of each device individually. Also, when adding a new device in Managed Devices, if they are in HA mode, both devices will be displayed together, so you can add both devices.			
	When viewing an HA pair, if the configuration does not match a warning indicator will appear. You will also see an indicator if the HA devices are in different device groups.			
	This option is also independent for each section, so enabling and disabling in one area, will not enable/disable for all areas. The Group HA Peers option is present in the following Panorama areas:			
	Managed Devices			
	• Templates			
	Device Groups			
	Policies tab (Target tab for all policy types)Commit dialog			

Table 148. Device Group Settings

Panorama Administrator Roles, Profiles, and Accounts

Panorama supports the following options to authenticate administrative users who attempt to log in to the device:

- Local database—The user login and password information is entered directly into the Panorama database.
- **RADIUS**—Existing Remote Authentication Dial In User Service (RADIUS) servers are used to authenticate users.

- LDAP—Existing Lightweight Directory Access Protocol (LDAP) servers are used to authenticate users.
- Kerberos—Existing Kerberos servers are used to authenticate users.
- Client Certificate—Existing client certificates are used to authenticate users.

When you create an administrative account, you specify client certificate (no authentication profile), or an authentication profile (RADIUS, LDAP, Kerberos, or local DB authentication). This setting determines how the administrator password is checked. If you do not specify a profile, the account will use local authentication.



Note: You may have certain Panorama administrators that may not have access to the **Panorama > Administrators** menu. In this case, the administrator can click on their username located to the left of the logout link on the bottom of the web interface to change their password.

Administrator roles determine the functions that the administrator is permitted to perform after logging in. You can assign roles directly to an administrator account, or define role profiles, which specify detailed privileges, and assign those to administrator accounts.

- Refer to the following sections for additional information:
- "Setting Up Authentication Profiles" on page 62.
- "Defining Panorama Administrator Roles" on page 388.
- "Certificate Profile" on page 89.
- "Specifying Panorama Access Domains for Administrators" on page 392.
- "Certificate Profile" on page 89.

Defining Panorama Administrator Roles

Panorama > Admin Roles

Use the **Admin Roles** page to define role profiles that determine the access and responsibilities available to administrative users. For instructions on adding administrator accounts, refer to "Creating Panorama Administrative Accounts" on page 389.



Note: The Admin Role can be mapped via RADIUS Vendor-Specific Attributes (VSA) using the following attribute: "PaloAlto-Panorama-Admin-Role = <AdminRoleName>,".

Table 149.	Panorama	Administrator	Role	Settings
------------	----------	---------------	------	----------

Field	Description
Name	Enter a name to identify this administrator role (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Description	Enter an optional description of the role.
Permission	Select the scope of administrative responsibility (Panorama or Device Group and Template).

Field	Description
WebUI	Click the icons for specified areas to indicate the type of access permitted
	for the web interface:
	• Read/write access to the indicated page.
	 Read only access to the indicated page.
	• No access to the indicated page.
XML API	Select the type of access for the XML API
	• Report —Access to the device reports.
	• Log—Access to the device logs.
	• Configuration —Permissions to retrieve or modify the device configura- tion.
	 Operational Requests—Permissions to run operational commands.
	• Commit—Permissions to commit the configuration.
	• User-ID Agent—Access to the User-ID Agent.
	• Export —Permissions to export files from the device, including the con- figuration, block or response pages, certificates, keys, and more.
	• Import —Permissions to import files to the device, including software, content, license, configuration, certificates, block pages, custom logs, and more.
Command Line	Select the type of role for CLI access:
	• None—Access to the device CLI not permitted.
	• superuser —Full access to the current device.
	• superreader —Read-only access to the current device.
	 panorama-admin—Full access to a selected device, except for defining new accounts or virtual systems.

Table 149. Panorama Administrator Role Settings (Continued)

Creating Panorama Administrative Accounts

Panorama > Administrators

Administrator accounts control access to Panorama. Each administrator can have full or readonly access to Panorama and all managed devices, or can be assigned Panorama administrator access, but will not have permissions to create admin accounts or edit admin roles, which allows access to Panorama configuration and not the managed devices. The predefined **admin** account has full access to Panorama and the managed devices.

The following authentication options are supported:

- Password authentication—The user enters a user name and password to log in. No certificates are required.
- Client certificate authentication (web)—If you select this check box, a user name and password are not required; the certificate is sufficient to authenticate access to the firewall.
- Public key authentication (SSH)—The user can generate a public/private key pair on the machine that requires access to the firewall, and then upload the public key to the firewall to allow secure access without requiring the user to enter a user name and password.



Note: To ensure that the device management interface remains secure, it is recommended that administrative change their passwords periodically using a mixture of lower-case letters, upper-case letters, and numbers. You can also enforce "Minimum Password Complexity" from Setup > Management.

Table 150. Administrator Account Settings

Field	Description
Name	Enter a login name for the user (up to 15 characters). The name is case-sensitive and must be unique. Use only letters, numbers, hyphens, and underscores.
Authentication Profile	Select an authentication profile for administrator authentication according to the settings in the specified authentication profile. This setting can be used for RADIUS, LDAP, or Kerberos authentication.
	For instructions on setting up authentication profiles, refer to "Setting Up Authentication Profiles" on page 62.
Use only client certificate authentication (Web)	Select the check box to use client certificate authentication for web access. If you select this check box, a user name and password are not required; the certificate is sufficient to authenticate access Panorama.
Password/Confirm Password	Enter and confirm a case-sensitive password for the user (up to 15 characters). You can also enforce "Minimum Password Complexity" from Setup > Management.
	You may have certain Panorama administrators that do not have access to the Panorama > Administrators menu, in this case, the administrator can click on their username located to the left of the logout link on the bottom of the web interface and change their local password from there.

Field	Description				
Use Public Key Authentication (SSH)	Select the check box to use SSH public key authentication. Click Import Key and browse to select the public key file. The uploaded key is displayed in the read-only text area.				
	Supported key file formats are IETF SECSH and OpenSSH. Supported key algorithms are DSA (1024 bits) and RSA (768-4096 bits).				
	<i>Note:</i> If the public key authentication fails, a login and password prompt is presented to the user.				
Role	Select an option for assigning a role to this user. The role determines what the user can view and modify.				
	• Dynamic , you can select any of the following pre-specified roles from the drop-down list:				
	- Superuser —Full access to the current device.				
	 Superuser (Read Only)—Read-only access to the current device. 				
	 Panorama administrator—Full access to the Panorama instances. 				
	• Role Based — Access based on assigned roles, as defined in "Defining Panorama Administrator Roles" on page 388.				
	If you choose Role Based, select a previously-defined role profile from the drop-down list. For instructions on defining role profiles, refer to "Defining Panorama Administrator Roles" on page 388				
	For role based access, when you select a profile assigned to the Device Group and Template administrator role, the Access Control tab is displayed. On this Access Control tab, you can define access to Device Groups, Templates, and Device Context. The definitions for these fields are the same as the information in "Specifying Panorama Access Domains for Administrators" on page 392.				

Table 150. Administrator Account Settings (Continued)



Note: On the Panorama **Administrators** page for "," a lock icon is shown in the right column if an account is locked out. The administrator can click the icon to unlock the account.

Specifying Panorama Access Domains for Administrators

Panorama > Access Domain

Use the **Access Domain** page to specify domains for role-based administrators who have access to device groups and templates. Adding a device group to an access domain allows you to manage policies and objects for that device group. Adding an individual firewall to an access domain allows you to switch into the device context for that firewall.

The access domain is linked to RADIUS vendor-specific attributes (VSAs) and is supported only if a RADIUS server is used for administrator authentication. If RADIUS is not used, the access domain settings on this page are ignored.



Note: The Access Domain can be mapped via RADIUS VSA using the following attribute: "PaloAlto-Panorama-Admin-Access-Domain = <AccessDomainName>,".

Field	Description
Name	Enter a name for the access domain (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, hyphens, and underscores.
Device Groups	Click Add to specify pre-defined device groups to include in the access domain.
Device Context	Select the device(s) that the administrator can do a context switch to in order to allow local configuration edits.
Templates	Click Add to specify pre-defined templates to include in the access domain.

Table 151. Access Domain Settings

Device Groups

Panorama device groups allow you to group firewall devices and then define policies and objects that can be shared across those device groups.

The following sections describe how to define policies and objects for device groups:

- "Working with Policies" in the next section
- "Working with Objects" on page 395

Working with Policies

Policies

Panorama allows you to define policies that are shared across the managed firewalls. You can apply pre and post rules that will apply to a device group and you can add additional global pre and post rules that apply to all device groups. This creates a layered approach of applying policies to managed devices. The first layer is the device level rules that are local to the device, you can then apply pre and post rules to device groups, and then you can add another layer of global pre and post rules that apply to all device group in the Panorama instance as shown in Figure 47.





General information about working with policies is found in "Policies" on page 183. This section describes the modifications and best practices that apply to policies in Panorama.

The following best practices apply to policies in Panorama:

• **Pre rules:** Pre rules are evaluated before any device specific rules and generally make up the majority of a deployment's shared rulebase. Do not add any pre rules if you will need device-level exceptions.

If you do not want administrators to be able to allow any applications at specific sites, you can include a deny rule for all zones, users, and applications as your last rule in the set of pre rules.

- **Firewall-specific rules:** Define rules for an individual firewall to create site-specific policies.
- **Post rules**: Use these rules to specify what happens to traffic that is not covered by the pre rules and firewall-specific rules. For example, if a pre rule specifies certain allowed applications and the post rule is "deny all," then applications not covered by the pre rule are stopped. You can then add rules to allow additional applications based on user request. You can also create device level "allow" rules as exceptions for specific applications that are permitted at an individual location.
- **Global pre rules:** Global pre rules are evaluated before any firewall-specific, or device group rules and are applied to all devices in managed device groups.
- **Global post rules:** Use these rules to specify what happens to traffic that is not covered by device group pre rules and firewall-specific rules. These rules are evaluated last, and only after the device group pre rules.

Global shared rules can only be created and modified by the Panorama Admin or superuser. So these rules can be used to apply policies before and after any rules applied by device group administrators.

The following applies when defining policies in Panorama:

Panorama applies policies to specified device groups and global shared policies applies another layer of policies to the device groups.

When you create a device group rule from the Panorama instance, you select the **Policies** tab, choose the device group that you want to create rules for and then define your rules. To create global rules that apply to all managed devices in device groups, you choose the **Shared** option in the **Device Group** drop down, as shown in Figure 48.

m naloalto					201	DIC	175	SIDAS	
Patoatto		Dashboard	ACC	Monitor	Policies	Objects	s Network	Device	Panorama
Context						6 m	10 L		10 II
Panorama		Device Group Sha	red						
🖯 📟 Security	× 4			1.55					
Pre Rules	11								
Post Rules	Nar	ne	Locat	ion Ta	g	Zone	Address	User	
Pre Rules									

Figure 48. Panorama Shared Policy

You can also specify that shared objects take precedence over device group objects by selecting the check box **Shared Objects Take Precedence** in **Panorama > Setup > Management > Panorama Settings**. This option is a system-wide setting and is off by default. When this option is off, device groups override corresponding objects of the same name. If the option is on (checked), device group objects cannot override corresponding objects of the same name from a shared location and any device group object with the same name as a shared object will be discarded.

• You can target a policy rule to individual devices within the device group for which the rule is defined. To target a device after a policy is created, click an entry in the **Target** column and select the devices in the pop-up window. If you do not select devices to target, the policy is added to all of the (unchecked) devices in the device group.

To apply the rule to all devices in a device group EXCEPT the targeted device, select the **Install on all but specified devices** check box and select the devices that you want to exclude. If you do not select any device, the policy is added to none of the devices in the device group.



Figure 49. Targeting Policy Rules to Individual Devices in Panorama

• Zones are not created in Panorama, but you can select zones based on zone names that are gathered from any template which maps to the same devices as the ones included in a device group. If a zone is not available from the templates, you must manually type a zone name when you first create a rule. For subsequent rules, you can enter new zones or select from previously entered zones.



Note: The gathering of objects from templates not only applies to zones, it applies to all objects that can be referenced from the Policy/Objects tab, which comes from the Device/Network tabs. This includes: zones, interfaces, certificates, Network Profiles > Monitor, server profiles (Syslog, Email, SNMP Traps), virtual systems in PBF rules, and the local user databases for users/groups.

• Each policy type listed on the side menu includes pages to define pre rules and post rules as well as global shared rules that are applied on top of the device group pre and post rules. See page 393 for information on best practices when using policies.

Working with Objects

Panorama supports sharing of objects defined in Panorama. You can create objects on Panorama and then push the object configurations to the managed firewalls. The objects become available for use in policies that are defined on the individual managed firewalls.



Note: All custom objects should have unique names, and predefined names such as "any" or "default" should be avoided. In particular, using the same object names with different device groups can cause confusion on devices and on Panorama.

All objects in the **Objects** tab and some objects in the **Device** tab can be managed centrally. **Device** tab objects are managed under the **Panorama** tab and include the following: certificates, response pages, server profiles (SNMP trap, syslog, email, RADIUS, LDAP, and Kerberos), authentication profiles and sequences, and certificate profiles. These objects have a **Location** field that allows you to select where the object should exist in the deployment (for example, "device-group-test"). The following table explains the available object assignment and sharing options for the **Location** field.

Field	Description						
Panorama	Panorama supports keeping objects locally and not pushing the objects to any managed devices. To do this, choose Panorama from the Location drop-down list when defining the object.						
	This option is available only on the Panorama tab and only for the following:						
	• Server profiles, including SNMP Trap, Syslog, Email, Remote Authenti- cation Dial In User Service (RADIUS), Lightweight Directory Access Protocol (LDAP), and Kerberos						
	• Authentication profiles, authentication sequences, and certificate pro- files						
Device Groups	Device groups are used to make objects and policies defined in Panorama available to specified sets of devices. For information on creating device groups, see "Defining Device Groups" on page 386.						
	 On the Policies or Objects tab, choose the device group from the Device Groups drop-down list when defining the object. 						
	Note: If you have objects of the same name where one is shared and another is device group specific, the device group specific object will be used for that device group.						
	Device Group dg1						
	• On the Panorama tab, choose a device group from the Location drop- down list when defining the object.						
Shared	Creating a shared object makes the object available for use in any device group. Only Panorama and administrators can create objects in the shared location.						
	• On the Panorama tab, choose Shared from the Location drop-down list when defining the object.						
	• On the Objects tab, select the Shared check box when defining the object.						

Table 152. Object Assignment and Sharing Options
Working with Devices

Switching context allows an administrator to switch from managing shared policy on Panorama to managing device-specific settings on an individual firewall (such as device specific policy, networking, and device setup). Use the **Context** drop-down list above the side menu to choose an individual device or the full Panorama view. You can select the name of any device that has been added for management by Panorama (refer to "Panorama Administrator Roles, Profiles, and Accounts" on page 387). When you select a device, the web interface refreshes to show all the device tabs and options, allowing you to manage all aspects of the device from Panorama.



Note: You can only switch context to connected devices. Disconnected devices are not shown in the drop-down list.

Context	Panorama	շխ
	Panorama	Y
	FW 102 FW 103	
	FVV 104	

Figure 50. Choosing Context

Commit Operation in Panorama

To commit Panorama configuration changes click the **Commit** icon to bring up the commit dialog box. This dialog box allows you to commit specific areas of the Panorama environment, see Figure 51.



Figure 51. Panorama Commit Dialog Box

The following options are available in the commit dialog box:



Note: When performing commit operations in Panorama, it is suggested that you commit the Panorama configuration first, before committing configuration updates to managed devices.

- **Commit Type**—Choose the commit type:
 - > **Panorama**—Commit the current candidate configuration for Panorama.
 - > **Template**—Commit template changes from Panorama to the selected devices. When committing templates, you can select a subset of devices if desired.
 - Device Group—Commit device configuration changes from Panorama to the selected device/virtual system(s).
 - Collector Groups—Only commit changes to Panorama log collector groups. This will commit changes made in the Panorama > Collector Groups page and will apply those changes to the log collector device.
- Include Device and Network Templates—This option is available when committing a Device Group from Panorama and is a combo operation that will include both the device and network template changes. The template that will be applied to the device is the template that the device belongs to as defined in Panorama > Templates. You can also select Commit Type Template to commit templates to devices.
- Force Template Values—When doing a Commit Type Template, you can select this option to remove objects on the selected devices or virtual systems that have been overridden by the local configuration. When doing a Commit Type "Device Group", you need to also select the "Include Device and Network Templates" check box since overriding can only occur for Template pushed configuration options. This will cause the overridden objects to inherit settings from the template. Refer to "Overriding Template Settings" on page 402.
- Merge with Candidate Config—Choose this option to cause the device to include its local candidate configuration when the commit is invoked from Panorama. If this option is not checked, the device local candidate config is not included.

It is important to leave this option unchecked when you have local administrators making changes on a device and you don't want to include their changes when pushing a configuration from Panorama.

Preview Changes—Click this button to bring up a configuration audit window that shows proposed changes in the candidate configuration compared to the current running configuration. You can choose the number of lines of context to display, or show all lines based on items that have been added, modified, or deleted. This option is available when using Commit Type Device Group, Template, or Panorama. You can also view the commit state for devices by navigating to Panorama > Managed Devices and viewing the Last Commit State column. The Device > Config Audit feature performs the same function, refer to "Comparing"

Configuration Files" on page 49.

After the commit is complete, you will see a "Commit succeeded" messages, if there are warning messages, you will see "Commit succeeded with warnings". To view warnings, navigate to **Panorama > Managed Devices** and see the **Last Commit State** column and click the text to view details.

Panorama Backward Compatibility

When upgrading Panorama to 5.0, and you are managing devices with PAN-OS 4.0 or earlier, simple style vulnerability protection and anti-spyware profiles are automatically converted to rules of equivalent meaning on those devices. Custom style profiles are converted to exceptions that specify signature-specific actions, with no rules required. After migration, a limited set of changes can be made to the migrated profiles in Panorama if compatibility with devices running PAN-OS 4.0 and earlier is required.

For rules created during conversion from a simple style profile, the action of the migrated rules can be modified, and additional allow exceptions can be added to the exceptions list. If a custom style profile was converted to an exceptions-based profile, the exceptions list can be freely modified, but no rules can be created. If the administrator attempts to commit using an incompatible profile, the commit will fail and the failure will be noted in the Managed Devices list under the Last Commit All State column.

To manage devices running different versions of software you must select, in the Device Groups configuration, a master device which is running the lowest software version of the User-ID agent to remain backward compatible and be able to push user based rules to older firewalls.

Templates

Panorama > Templates

The Panorama **Templates** page is used to create templates that can be used to manage configuration options based on the Device and Network tabs, enabling you to deploy configurations to multiple devices that require similar settings. You can also deploy a base configuration and if needed, override specific settings on the device. For example, you can deploy a base config to a global group of devices, but configure specific time zones settings directly on the devices based on their location.

When managing device configuration with Panorama, you can use a combination of **Device Group** configuration (to manage shared policies and objects) settings and **Templates** settings (to apply device and network settings), but these features are managed separately because of the differences in what can be configured. For information on adding and configuring Panorama templates, refer to "Configuring Panorama Templates" on page 401.

Field	Description
Name	Enter a template name (up to 31 characters). Use only letters, numbers, spaces, hyphens, periods, and underscores. The name is case-sensitive and must be unique.
	This name will appear in the Device and Network tab in the Template drop- down menu. When selecting a template from one of these tabs, the settings that are modified will only apply to the selected template.
Description	Enter a description for the template.
Virtual Systems	Select the check box if the template will be used on devices with multiple virtual systems. When defining template settings for multi-virtual system devices, you need to configure settings for each virtual system on the device. Note: A template enabled for devices with multi-virtual systems (multi-vsys)
Note: You cannot use a template to create virtual systems on the device.	cannot be pushed to devices with a single virtual system. When you upgrade your Panorama server to v5.0 and later, by default templates are created for configurations relating to the network and device tabs. If, for example, you have defined a server profile for the managed firewall(s), a template is automatically generated for the server profile. This auto-generated template is enabled for multiple virtual systems. To prevent a commit failure, make sure to clear the Virtual Systems checkbox before you push the template to a device that is not multi-vsys capable or if the multi- vsys capability is disabled on the device.
Operational Mode	Set the PAN-OS operational mode for the template: normal, fips, or cc.
VPN Disable Mode	Selecting this check box will hide all VPN related options in the Device and Network tabs. The ability to install GlobalProtect Portal or Gateway licenses is also disabled in this mode.
	Note: This option is designed for countries that do not allow VPN connectivity. Palo Alto Networks hardware models that have the -NV indicator in the model name are hard coded to not allow VPN configurations, so this option should be used when creating templates for these models.

Table 153 Template Settings (Panorama)

Description
This window lists all devices managed by Panorama and will also show device groups. Click the check box next to devices to make them a member of the template. You can also add all devices in a device group by clicking the check box next to the device group.
Note: Templates are based on individual devices, not device groups. Adding a new device to a device group will not automatically add it to a template.
Select the check box to group devices in high availability (HA) mode together.
This option allows you to easily identify devices that are in HA mode. When pushing shared policies, you can push to the grouped pair, instead of each device individually. Also, when adding a new device in Managed Devices, if they are in HA mode, both devices will be displayed together, so you can add both devices.
When viewing an HA pair, if the configuration does not match a warning indicator will appear. You will also see an indicator if the HA devices are in different device groups.
This option is also independent for each section, so enabling and disabling in one area, will not enable/disable for all areas. The Group HA Peers option is present in the following Panorama areas:
Managed Devices
• Templates
Device Groups
 Policies tab (larget tab for all policy types) Commit dialog

Table 153 Template Settings (Panorama)

Configuring Panorama Templates

Panorama > Templates

To configure Panorama templates, you first create the template and then add devices to it. After the first template is created, you will then see a **Template** drop-down menu in the **Device** and **Network** tabs. Select the desired template from the **Template** drop-down menu and configure device and network settings as if you were managing a single device, but all options that are set will only be applied to the selected template. After the template is configured, you can do a commit from Panorama that only applies to templates.

To create and configure a template, follow these steps:

Adding a New Template

- 1. Under the **Panorama** tab, click **Templates** to open the templates page.
- 2. Click **Add** and enter the template configuration options. Refer to "Template Settings (Panorama)" on page 400.
- 3. In the **Devices** field, you will see a list of all devices that are managed by Panorama. Click the check box next to each item to make them a member of the new template. If you select a device group, all devices in that group will be selected.
- 4. Click **OK** to save the template.

Configuring a Template

1. Now that a template has been created, click the **Device** or **Network** tab and you will see a **Template** drop-down menu, as shown in Figure 52.

Figure 52. Template Menu



- 2. Click the **Template** drop-down and select the template that you want to configure.
- 3. Click the **Device** or **Network** tab and set the desired configuration options for the template.



Note: When a template is selected from the Device or Network tabs, you can only set options that are part of the configuration. You can not set operational type options such as changing the device to multi-virtual system mode, and setting a Master Key.

- 4. After making all of the configuration changes, click **Commit** and from the **Commit Type** drop-down menu select **Template**. You can also use the **Device Group** commit option and select **Include Device and Network Templates** check box to push templates to a device group.
- Click the check box next to each template that you want to commit and then click OK. You can also preview your changes from the Commit window by clicking the Preview Changes button. A pop-up will appear showing commit status.

Overriding Template Settings

When you apply a template to control device and network settings on a firewall, you may want to override some of those settings and have them controlled by the local device configuration. Example, you can deploy a base config to a global group of devices, but configure specific time zones settings directly on the devices based on their location using an override.

To override device and network setting applied by a template, you simply change to the device context, or access the device directly, navigate to the desired setting and then click the **Override** button. The setting will be copied to the local configuration of the device and will no longer be controlled by the template. You can also revert the change by clicking the **Restore** button and the setting will once again be inherited from the template. When doing a commit from Panorama to a managed device that contains overrides, you can select the **Force Template Values** check box to have Panorama templates take over any overridden objects.

When overriding **Device > Setup and Device > High Availability** settings, the overrides are for individual values and parameters inside of configuration trees, and are not applied to an entire tree configuration. This includes items such as DNS servers, Management IP, or NTP server settings. For items such as interfaces and RADIUS server profiles, you apply overrides to the entire object, not internal values.

To identify settings that have templates applied, you will see the following indicators as shown in Figure 53:

Figure 53. Template Indicators



The single green icon indicates that a templates has been applied and there are no overrides. The green and orange icon indicates that a template has been applied and some settings have been overridden.

Removing Templates

To remove a template, you must disable the template on the local device. On the managed device, navigate to **Device > Setup > Management** tab, then edit the **Panorama Settings** page and click the **Disable Device and Network Template** button. Removing the device from the configuration in **Panorama > Templates** will not delete the template values on the local device.

Logging

Panorama performs two functions: device management and log collection. To facilitate scalability in large deployments, you can use the M-100 appliance to separate the management and log collection functions on Panorama.

The following sections describe options available for log collection:

- "Logging and Reporting" in the next section
- "Using Panorama for Log Collection" on page 404

Logging and Reporting

The Panorama logs and reports provide information about user activity in the managed network. Report statistics are aggregated every 15 minutes for use in scheduled predefined and custom reports and statistics are forwarded to Panorama on an hourly basis. If log forwarding is enabled, logs are sent when they are generated on the device.

The **ACC** tab in Panorama displays information from the connected firewalls, all the tables pull information dynamically from the firewalls; they do not require explicit log forwarding. Log forwarding is required for long term log storage and for reporting on logs stored locally in Panorama.

Generating User Activity Reports

Monitor > PDF Reports > User Activity Report

The Panorama user activity report summarizes user activity across all of the managed firewalls. It is based on firewall data that has been forwarded to Panorama. Refer to "Managing User Activity Reports" on page 272 for general information on creating user activity reports.

Using Panorama for Log Collection

The M-100 appliance provides a comprehensive log collection solution for Palo Alto Networks firewalls. It helps offload the intensive log collection process from your Panorama management server. Once deployed, each firewall in your managed environment can be configured to send logs to one or more log collectors. By configuring multiple log collectors per firewall, storage capacity is spread out among the collectors creating a more flexible environment with redundancy.

The ACC, PDF Reports, and Logs viewer on Panorama are used to query aggregated information for all managed firewalls. This is the same process that is used if logs were sent directly to Panorama, but in this case, Panorama queries the log collectors to gather the information.

For information on the M-100 hardware, refer to the M-100 Hardware Reference Guide.



Note: The M-100 appliance is shipped with Panorama pre-loaded and both Panorama management and log collection functions are enabled by default. However, if you plan to use the M-100 for log collection only, you need to follow these steps:

- 1. Perform initial configuration to set up an IP address for the management interface and change the default password on the appliance.
- 2. Apply the license on the appliance.
- 3. Verify that the firewalls on your network are running PAN-OS 5.0. If not, you must first upgrade to you PAN-OS 5.0 because support for separate appliances for Panorama management and log collection is not available on earlier versions of PAN-OS
- 4. Change the operational mode on the appliance to log collector mode. For details on the CLI commands, see "Configuring the M-100 as a Log Collector" on page 405.
 When the M-100 is in log collector mode, there is no Web interface; Only CLI is available for managing/configuring the appliance.

The following sections describe how to deploy log collectors:

- "Deploying Distributed Log Collection" in the next section
- "Configuring the M-100 as a Log Collector" on page 405
- "Configuring the Panorama Server for Log Collector Management" on page 406
- "Managing Log Collectors" on page 407
- "Defining Log Collector Groups" on page 411
- "Log Collector Storage" on page 414

Deploying Distributed Log Collection

Several things should be considered when deploying a distributed log collection solution, including:

- Determine the location of the M-100 appliances based on your network topology and the location of your firewalls. Ideally, the Panorama management server, firewalls, and log collectors should all be connected over a management network. You can configure the firewalls to connect to multiple log collectors in case of failures, so its important that each firewall can reach its assigned collector(s). Typically, the Panorama server and log collectors would be installed in data centers.
- Determine the disk storage needs for log collection based on the number of firewalls that it will handle and the desired retention time. You should analyze your current environment to determine the amount of logs and reports that are being generated.

The illustration in Figure 54 shows a basic log collection deployment. Managed firewalls are configured to send log information to the M-100 log collector. The Panorama servers in HA mode then communicate with the log collectors to perform reporting and to view log information collected from the managed firewalls.

In large deployments where multiple log collectors are available, you assign two or more log collectors to each firewall in a preference order. The first log collector you specify will be the primary log collector for the firewall. If the primary log collector fails, the firewall will momentarily cache log information, so nothing is lost, and will then send logs to the secondary log collector. Regardless of firewall to log collector assignments, logs are balanced across all log collectors in the group to maintain even storage usage among the collector group.

Figure 54. Distributed Log Collection



Configuring the M-100 as a Log Collector

To configure the M-100 appliance as a log collector, follow these steps:

- 1. Perform initial configuration to set up an IP address for the management interface (MGT port) and change the default password on the appliance. The MGT port is used for all communications, including device management, log collection, and log collector to log collector communications.
- 2. Apply a license to the M-100.
- 3. Before you proceed, verify that the firewalls on your network are running PAN-OS 5.0. If not, you must first upgrade the firewalls to PAN-OS 5.0. This upgrade is necessary because support for separate appliances that perform the Panorama management and log collection functionality is not available on earlier versions of PAN-OS.
- 4. Convert the appliance to function as a log collector only. To change the functionality of the appliance to logger-mode:
 - a. Log in to the CLI. For instructions on accessing the CLI, refer to the *Command Line Interface Reference Guide*.
 - b. Enter the following CLI command:

request system logger-mode logger

Answer Yes to confirm the change. The device will reboot.

c. After the appliance boots up, verify that the device is in logger mode, using the CLI command:

```
show system info | match logger_mode
```

The output should display: logger-mode: True



Note: When the M-100 is in log collector mode, only the CLI is available for management. Most of the configuration is performed on the Panorama manager.

5. Configure the IP address or FQDN of the Panorama server that will be used to manage the log collector.

set deviceconfig system panorama-server ip-address <Panorama-IP>. If your Panorama server is in HA mode, enter the IP address of the peer device: set deviceconfig system panorama-server-2 ip-address <Panorama-IP>.

6. Type **commit** to make the change active and then **exit** to leave configuration mode.

Now that the log collector has the basic configuration needed, for further configuration on the log-collector, use the Panorama Management server to complete the configuration. Refer to "Configuring the Panorama Server for Log Collector Management" in the next section.

Configuring the Panorama Server for Log Collector Management

The Panorama management server is used to configure and manage the M-100 log collector once network connectivity is established between the two systems. From Panorama, you manage all of the log collector settings, storage, log retention, SNMP settings, and the assignment of managed firewall to log collectors using collector groups. Once the log collectors are in place, Panorama then queries those log collectors for aggregated reporting. Refer to "Defining Log Collector Groups" on page 411.



Note: You cannot migrate existing firewall logs to the M-100; only new log data will be forwarded after the configuration is completed.

For detailed descriptions of each of the Managed Collectors fields, refer to "Managing Log Collectors" on page 407.

To configure log collectors from Panorama:

- 1. Make sure the log collector you want to manage has basic network connectivity so the Panorama management server can communicate with it. Refer to "Configuring the M-100 as a Log Collector" on page 405.
- 2. From the Panorama management server, navigate to **Panorama > Managed Collectors** and then click **Add** to start adding a log collector.
- 3. In the **General** tab, specify the log collector serial number and log collector name (hostname). Then specify the Panorama server IP address, DNS, NTP, Timezone, and location information.
- 4. Use the **Authentication** tab to update the local admin account password on the log collector. You can only use a password hash for this field. You can create a password hash using the Panorama CLI command **request password-hash password** *password*. Press enter and the hash value will be displayed. Copy the hash and paste into the **Password Hash** field and **Confirm Password hash** field.
- 5. The **Management** tab is used to configure the management port on the log collector and is labeled MGT on the front of the M-100. You can set the MTU and other interface settings, the allowed services and define a list of permitted IP addresses that will be allowed to manage the log collector. The MGT port will be used for all communications between the log collector and managed devices. Refer to "Defining Log Collector Groups" on page 411.
- 6. If you are increasing storage capacity, see "Log Collector Storage" on page 414 for information on configuring the additional disks as a RAID pair.
- 7. Click **Commit** and in the **Commit Type** drop down select **Panorama** and click **OK**.
- 8. Click **Commit** and in the **Commit Type** drop down select **Collector Group**, select the log collector or group that you would like to commit and then click **OK**. This will push the changes to the selected log collector(s). You will see a pop-up window that will show the commit state.
- 9. The log collector should now be configured. The next procedure is to create **Collector Groups** to assign firewalls to log collectors. Refer to "Defining Log Collector Groups" on page 411.

Managing Log Collectors

Panorama > Managed Collectors

Use the Managed Collectors page to configure, manage, and update log collector devices. The settings on this page also exist in the log collector CLI.

After you add log collectors, you can click the **Statistics** link for each collector. This will show the **Collector Statistics** window where you can view Disk Information, Performance numbers for the CPU and Average Logs/sec. You can also view information on the Oldest log received by the collector to get a better understanding of the log range you are looking at.

Field	Description
General Tab	
Collector S/N	Enter the serial number of the log collector device.
Collector Name	Enter a name to identify this log collector (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
	This is the hostname of the log collector.
Panorama Server IP	Specify the IP address of the Panorama server used to manage this collector.
Panorama Server IP 2	Specify the IP address of the secondary device if the Panorama management server is in HA mode.
Primary DNS Server	Enter the IP address of the primary DNS server. The server is used for DNS queries from the log collector, for example, to find the Panorama server.
Secondary DNS Server	Enter the IP address of a secondary DNS server to use if the primary server is unavailable (optional).
Primary NTP Server	Enter the IP address or host name of the primary NTP server, if any. If you do not use NTP servers, you can set the device time manually.
Secondary NTP Server	Enter the IP address or host name of secondary NTP servers to use if the primary server is unavailable (optional).
Timezone	Select the time zone of the log collector.
Latitude	Enter the latitude (-90.0 to 90.0) of the log collector that is used in the traffic and threat maps for App-Scope.
Longitude	Enter the longitude (-180.0 to 180.0) of the log collector that is used in the traffic and threat maps for App-Scope.
Authentication Tab	
User Name	This field will always show admin and is used for the local CLI login name on the log collector.

 Table 154
 Managed Collectors Page

Field	Description
Mode	Select Password to manually enter and confirm a password to be used for local CLI authentication, or select Password Hash to enter a hash value.
	To create a password hash from the Panorama management server CLI, run the following:
	request password-hash password password123
	This will return a hash value for the password <i>password</i> 123 (<i>Example-</i> \$1\$ <i>urlishri</i> \$ <i>a</i> LP2 <i>by.u</i> 2 <i>A</i> 11 <i>Q</i> / <i>Njh</i> 5 <i>T</i> F <i>y</i> 9).
	Copy the hash value from the CLI and paste to the Password Hash field. When you commit your changes, the new hash will be pushed to the log collector and the new local admin login password will be <i>password123</i> .
Failed Attempts	Specify the number of failed login attempts that are allowed for the web interface and CLI before the account is locked. (1-10, default 0). 0 means that there is no limit.
Lockout Time (min)	Specify the number of minutes that a user is locked out (0-60 minutes) if the number of failed attempts is reached. The default 0 means that there is no limit to the number of attempts.
Management Tab	Configure the management port settings labeled MGT on the front of the device. This port is used for all log collector communication.
Interface	The interface cannot be changed and the default label is MGT.
Speed and Duplex	Select the interface speed in Mbps (10, 100, or 1000) and the interface transmission mode full-duplex (Full), half-duplex (Half), or negotiated automatically (Auto).
IP Address	Enter the IP Address of the log collector management interface. The default IP address is 192.168.1.1.
Netmask	Enter the network mask for the IP address, such as "255.255.255.0".
Default Gateway	Enter the IP address of the default router (must be on the same subnet as the management port).
IPv6 Address	Enter the IPv6 address of the log collector management interface.
IPv6 Default Gateway	Enter the IPv6 address of the default router (must be on the same subnet as the management port).
MTU	Enter the maximum transmission unit (MTU) in bytes for packets sent on this interface (range 512 to 1500, default 1500).
Management Interface Services	Select the services that are enabled on the managed interface on the log collector device:
	• SSH—Select the check box to enable secure shell.
	 Ping—Select the check box to enable ping.
	• SNMP—Select the check box to enable the Simple Network Managed Protocol.

Table 154 Managed Collectors Page

Field	Description
Permitted IP Addresses	Click Add to enter the list of IP addresses from which management is allowed.
Disks Tab	Click Add to define the RAID 1 disk pair that will be used to store logs. You can then add additional disk pairs as needed to expand your storage capacity.
	By default, the M-100 is shipped with the first RAID 1 pair enabled with drives installed in bays A1/A2. You can add up to 3 more RAID 1 pairs to increase storage capacity by putting in new RAID 1 pairs in bays B1/B2, C1/C2, and D1/D2. In the software, the RAID 1 pair in bays A1/A2 is named Drive Pair <a b c d>.</a b c d>
	Refer to "Log Collector Storage" on page 414.

 Table 154
 Managed Collectors Page

Defining Log Collector Groups

Panorama > Collector Groups

Collector groups are used to assign Panorama managed firewalls to log collectors that will be used to offload the work of log collection that would normally be handled by the Panorama management server. Once the log collectors are in place and the firewalls are configured, the defined logs for each device will be sent to the log collectors and Panorama will then query the log collectors for aggregated log viewing or investigation. You also use collector groups to define storage retention and SNMP settings.

For detailed descriptions of each of the Collector Groups settings, refer to "Collector Groups Settings" on page 412.

To configure log collector groups:

- 1. From the Panorama management server, navigate to **Panorama > Collector Groups**.
- 2. Click **Add** and enter a name to identify the collector group.
- 3. In the **Min Retention Period (days)** field, enter the number of days that should be maintained across all log collectors in the group before an alert is generated. An alert violation in the form of a system log will be generated if the current date minus the oldest log is less than the defined min retention period.
- 4. Click the **Monitoring** tab and enter the SNMP settings for your environment in order to manage your devices through a system management solution.
- 5. Click the **Log Forwarding** tab and in the **Collectors** window, click **Add** and select the log collectors that will be part of this collector group. This list is generated by the collectors defined in **Panorama > Managed Collectors**.



Note: When you add the device serial number to the collector group, the managed device will start to send all logs to the collector group. To have the managed device revert back to sending logs to the Panorama manager, just remove the device from the collector group. This would also be required when migrating managed devices to a different install of Panorama manager. If you forget to do this, you will need to run the operational command delete log-collector preference-list from the managed device.

- 6. Configure the mapping that defines which devices forward to a preferred list of collectors in preference order. In the **Log Forwarding** tab, click **Add** under **Devices** and **Collectors** list windows and the **Devices** window will appear.
- 7. In the **Devices** drop-down, select a managed firewall and then under **Collectors** click **Add** and select a log collector. If you select multiple collectors the first collector added will be the primary collector, the second will be the secondary collector, and so on. If the primary fails, logs will be sent to the secondary. Click **OK** to save your changes.
- 8. Click **Commit** and in the **Commit Type** drop down select **Panorama** and click **OK**.
- 9. Click **Commit** and in the **Commit Type** drop down select **Collector Group** and click **OK**. Once the collector group commit completes, the managed firewalls that have been assigned to collectors will start forwarding logs to the collectors.

Now that firewall logs are being sent to the log collectors, you then use Panorama's ACC, PDF Reports, and the Logs viewer to query aggregated information for all of your managed firewalls. The ACC will query data directly from the log collectors, which

means that the ACC is looking at forwarded data from the firewall. If you deploy an M-100 as a Panorama manager and log collector, the ACC will also query the data forwarded to the log collector that is on the same device in this case.

Field	Description
General Tab	
Name	Enter a name to identify this collector group name that will be used to group log collectors for configuration and software update purposes (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
Log Storage	Indicates the current log storage quotas for the collector group. If you click on the capacity text, the Log Storage Settings window will appear. From here, you can allocate storage to various log features, such as Traffic, Threat, Config, System, and Alarm. You can also click Restore Defaults to use the default log allocation settings.
Min Retention Period (days)	Specify the retention period in days that should be maintained across all log collectors in the group before an alert is generated. An alert violation in the form of a system log will be generated if the current date minus the oldest log is less than the defined min retention period (range 1-2000 days).
Monitoring Tab	

Table 155 Collector Groups Settings

Field	Description
SNMP	The SNMP option enables you to collect information about the log collectors, including: connection status, Disk drive statistics, software version, average CPU, Average log/second, and storage duration per DB type (e.g. minutes, hours, days, weeks). SNMP information is based on a per-collector group.
	Specify the SNMP settings:
	 Location—Specify the location of the log collector device.
	• Contact—Specify an email contact for this device.
	• Access Setting—Specify the SNMP version that will be used to communicate with the Panorama management server (V2c or V3).
	If you select V3, specify the following:
	 Views— Click Add and configure the following settings:
	 View—Specity a name for a view.
	> OID—Specify the object identifier (OID).
	 Option (include or exclude)—Choose whether the OID is to be included or excluded from the view.
	 Mask—Specify a mask value for a filter on the OID in hexadecimal format (for example, 0xf0)
	 Users—Click Add and configure the following settings:
	 Users—Specify a user name that will be used for authentication between the log collector and SNMP management server.
	 View—Specify the group of views for the user.
	 Authpwd—Specify the user's authentication password (minimum 8 characters). Only Secure Hash Algorithm (SHA) is supported.
	 Privpwd—Specify the user's encryption password (minimum 8 characters). Only Advanced Encryption Standard (AES) is supported.
	 SNMP Community—Specify the SNMP community string that is used by your SNMP management environment. SNMPv2c Only (default is public).

Table 155 Collector Groups Settings (Continued)

Field	Description
Collectors	Click Add and select the log collector from the drop-down that will be part of this group. The drop-down will show all log collectors that are available in the Panorama > Managed Collectors page.
Devices	Click Add and then click the Devices drop-down and select the managed firewall that will be part of this collector group.
	Click Add in the Collectors window and select the collector that you would like to assign this firewall for log forwarding.
	Click OK to save your changes.
	When viewing the Devices window, the Devices column will list each firewall and the Collectors column will list the assigned collector(s) for the firewall.
	The first collector you specify will be the primary collector for the firewall. If the primary collector fails, the firewall will then send logs to the secondary collector. If the secondary fails, then the tertiary collector will be used, and so on.
	Note: When you add the device serial number to the collector group, the managed device will start to send all logs to the collector group. To have the managed device revert back to sending logs to the Panorama manager, just remove the device from the collector group. This would also be required when migrating managed devices to a different install of Panorama manager. If you forget to do this, you will need to run the operational command delete log-collector preference-list from the managed device.

Table 155 Collector Groups Settings (Continued)

Log Collector Storage

As your environment grows, you will need to increase the storage capacity of the log collector M-100 appliances, or you may need to add new log collectors.

By default, the M-100 is shipped with the first RAID 1 pair enabled with drives installed in bays A1/A2. You can add up to 3 more RAID 1 pairs to increase storage capacity by adding in new RAID 1 pairs in bays B1/B2, C1/C2, and D1/D2. In the software, the RAID 1 pair in bays A1/A2 is named Disk Pair A, B1/B2 is Disk Pair B, C1/C2 is Disk Pair C, D1/D2 is Disk Pair D.

For information on physically replacing drives on the M-100, refer to the M-100 Hardware *Reference Guide*.



Note: The RAID 1 disk arrays on the M-100 requires that each mirrored disk pair be installed in the appropriate disk bays. The device ships with A1/A2 drives installed and mirrored. To add more disks, you install the next pair into the next Disk Pair, which is bays B1/B2, then C1/ C2, and then D1/D2. You cannot install new disks in B1 and C1 for example, and then try to mirror those.

To expand disk storage on an M-100 using the B1/B2 drive bay:

- 1. Obtain two identical M-100 disk drives from Palo Alto Networks.
- 2. You do not have to power down the log collector to add new drives. If you prefer to power down the device, run **request shutdown system** from the CLI.

- 3. Remove the empty drives bays B1/B2 and install the new disks into the drive bays for Disk Pair B.
- 4. Insert the drive bays back into bays B1/B2.
- 5. Go to the CLI on the log collector and run **request system raid add B1**. This will initialize and format the B1 drive.
- 6. Run the command again on drive B2 **request system raid add B2**. The Raid 1 mirror will be created and the drives are now ready to be added to the collector group from the Panorama management server.



Note: The time taken to mirror the data on the drive may vary from several minutes to a couple hours, depending on the amount of data on the drive. Use the command show system raid detail to monitor the progress of the RAID configuration.

- Go to the Panorama management server that is managing this log collector, navigate to Panorama > Manage Collectors and click the log collector to open the Collector window. Select the Disks tab.
- 8. Disk A should already exist. Click **Add** and select **Disk Pair B** and then click **OK** to make the new disk pair available to the system.

When you increase or decrease storage on a log collector, the new storage capacity information is recognized by the Panorama management server and the log collector groups will be updated to re-balance logging. Regardless of firewall to collector assignments, logs are balanced across all collectors in the group to maintain even storage usage among the collector group. Note that during this re-balancing process, you may see heavy network traffic as log data is moved between collectors.

To view the disk capacity for a log collector group, or an individual log collector, go to the web interface on the Panorama management server, navigate to **Panorama > Collector Groups**, click a group name and then click the **General** tab. Next to **Log Storage**, you will see a link that displays the total and free storage information. If you click this link, the **Log Storage Settings** windows will appear where you can allocate storage for various features.

Viewing Firewall Deployment Information

Panorama > Device Deployment

Open the **Device Deployment** pages to view current deployment information on the managed devices and manage software versions on the devices, as described in the following table.

Field	Description
Software	Lists the versions of firewall software that are available for installation on the managed firewalls.
SSL VPN Client	Lists the versions of SSL VPN client software that are available for installation on the managed firewalls.
GlobalProtect Client	Lists the versions of GlobalProtect client software that are available for installation on the managed firewalls.
Dynamic Updates	Lists the threat and application definitions that are available for use on the managed firewalls. Refer to "Updating Threat and Application Definitions" on page 55 for information on using this page.
Licenses	Lists each managed device and the current license status. Each entry indicates whether the license is active (\bigcirc icon) or inactive (\bigotimes icon), along with the expiration date for active licenses.
	Perform either of the following functions on this page:
	• Click Refresh to update the list.
	• Click Activate to activate a license. Select the managed devices for activation and enter the authentication code that Palo Alto Networks provided for the device.

Table 156. Panorama Deployment Pages

Perform any of the following functions on the **Software**, **SSL VPN**, or **GlobalProtect** pages:

- Click **Refresh** to view the latest software releases available from Palo Alto Networks.
- Click **Release Notes** to view a description of the changes in a release.
- Click **Download** to install a new release from the download site. When the download is complete, a checkmark is displayed in the **Downloaded** column. To install a downloaded release, click **Install** next to the release.

During installation, you are asked whether to reboot when installation is complete. When the installation is complete, you will be logged out while the firewall is restarted. The firewall will be rebooted, if that option was selected.

- Click **Upload** to install or activate a release that you previously stored on your PC. Browse to select the software package, and click **Install from File**. Choose the file that you just selected from the drop-down list, and click **OK** to install the image.
- Click the **Delete** icon 🔀 to delete an outdated release.

Backing Up Firewall Configurations

Panorama > Setup

Panorama automatically saves every committed configured from the managed firewalls. You can configure the number of versions to keep on the Panorama device by using the Management settings under **Setup** on the **Panorama** tab. The default is 100. For instructions on configuring the number of versions, refer to "Defining Management Settings" on page 30.

To manage backups on Panorama, choose **Panorama > Managed Devices** and click **Manage** in the **Backups** column for a device. A window opens to show the saved and committed configurations for the device. Click a **Load** link to restore the backup to the candidate configuration, and then make any desired changes and click **Commit** to restore the loaded configuration to the device. To remove a saved configuration, click the 🕱 icon.

Scheduling Configuration Exports

Panorama > Scheduled Config Export

Panorama saves a backup of running configurations from all managed devices in addition to its own running configurations. Use the **Scheduled Config Export** page to collect the running configurations from all of the managed devices, package them in one gzip file, and schedule the package for daily delivery to an FTP server or by using Secure Copy (SCP) to transfer data securely to a remote host. The files are in XML format with file names that are based on the device serial numbers.

Use this page to set up a schedule for collection and export of the managed device configurations.

Field	Description
Name	Enter a name to identify the configuration bundle export job (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, hyphens, and underscores.
Description	Enter an optional description.
Enable	Select the check box to enable the export job.
Scheduled export start time (daily)	Specify the time of day to start the export (24 hour clock, format HH:MM).
Protocol	Select the protocol to use to export logs from the firewall to a remote host. You can use SCP to export logs securely, or you can use FTP, which is not a secure protocol.
Hostname	Enter the IP address or host name of the target FTP server.
Port	Enter the port number on the target server.
Path	Specify the path located on the FTP server that will be used to store the exported information.
Enable FTP Passive Mode	Select the check box to use FTP passive mode.
Username	Specify the user name on the target system.

 Table 157.
 Scheduling Configuration Bundle Exports

Field	Description
Password	Specify the password for the user on the target system.
Confirm Password	
Test SCP server connection	Click this button to test communication between the firewall and the SCP host/server.
	To enable the secure transfer of data, you must verify and accept the host key of the SCP server. The connection is not established until the host key is accepted.

 Table 157.
 Scheduling Configuration Bundle Exports (Continued)

Upgrading the Panorama Software

▶ Panorama > Software

To upgrade to a new release of Panorama software, you can view the latest versions of the Panorama software available from Palo Alto Networks, read the release notes for each version, and then select the release you want to download and install (a support license is required).

To upgrade the Panorama software, click **Refresh** to view the latest software releases available from Palo Alto Networks. To view a description of the changes in a release, click **Release Notes** next to the release.



Note: Panorama periodically performs a file system integrity check (FSCK) to prevent corruption of the Panorama system files. This check occurs after 8 reboots or at a reboot that occurs 90 days after the last file system integrity check was executed. If Panorama is running a FSCK, you will see a warning on the web interface and SSH login screens indicating that an FSCK is in progress and you cannot log in until it completes. The time to complete this process varies by the size of the storage system; depending on the size, it can take several hours before you can log back into Panorama.

To view progress, set up console access to Panorama.

- 1. To install a new release:
 - a. Click **Download** next to the release to be installed. When the download is complete, a checkmark is displayed in the **Downloaded** column.
 - b. To install a downloaded release, click **Install** next to the release.

When the installation is complete, you will be logged out while the Panorama system is restarted.

2. To delete an outdated release, click \mathbf{X} next to the release.



Note: Software is deleted to make room for newer version downloads. This happens automatically and cannot be manually controlled.

Upgrading the Panorama Software

Chapter 14 Configuring WildFire

This chapter describes how to use WildFire for analysis and reporting on malware that traverses the firewall:

- "About WildFire" in the next section
- "Setting Up WildFire on the Firewall" on page 423
- "Configuring WildFire Forwarding" on page 424
- "WildFire Data Filtering Log" on page 425
- "Using the WildFire Portal" on page 425

About WildFire

WildFire allows you to securely send files to the Palo Alto Networks cloud-based malware analysis center, where the files are automatically analyzed in a virtual sandbox for malicious activity. Files can be submitted automatically from firewalls through the use of a file blocking policy, manually through the WildFire Portal, or using the WildFire API (WildFire subscription required). The WildFire system executes the file in a virtual environment and watches for many potentially malicious behaviors and techniques, such as modifying critical system files, disabling security features, or using a variety of methods to evade detection. As new malware is detected, WildFire automatically generates antivirus signatures and distributes these signatures to firewalls with a threat prevention subscription via the daily antivirus signature updates. If you are a WildFire subscriber, signatures are available on a sub-hourly basis. Supported file types include Win32 Portable Executable (PE) files (e.g. exe, dll, and scr). When choosing file types in the objects profile, you can choose **PE** to cover all Win32 PE file types.

File types can be analyzed even if they are compressed (zip, gzip) or over SSL if decryption is enabled in the policy. You do not need to add (zip, gzip) to the file blocking profile to support analysis of compressed PE files since this is done automatically.



Note: When choosing **PE** in the objects profile **File Types** column to select a category of file types, do not add an individual file type that is part of that category or you will receive redundant entries in the **Data Filtering** logs. For example, if you choose **PE** you would not want to additionally include **exe**. This also applies to the **zip** file type, since supported file types that are zipped will automatically be sent to WildFire.

Choosing a category will also ensure that as new file type support is added to a given category, they will automatically be part of your object profile. You can also select **any** to cause all supported file types to be uploaded to WildFire.

As new malware is detected, WildFire will automatically generate antivirus signatures and will distribute these signatures to firewalls that have a threat prevention subscription. The signatures will be part of the daily antivirus signatures that are updated every 24-48 hours. If you have a WildFire subscription, you will have several new features enabled, including:

If you have a WildFire subscription, you will have several new features enabled, including:

- WildFire Dynamic Updates—A new WildFire section in Device > Dynamic Updates is now available. As new malware is detected, WildFire signatures are generated every 30 minutes, and the firewall's WildFire update schedule can be configured to poll for new WildFire signatures every 15, 30, or 60 minutes. As new WildFire signatures are downloaded to your firewall, you can configure the firewall to take specific actions on these signatures, separate from the regular antivirus signature actions in the antivirus profile.
- Integrated WildFire Logs—As files are uploaded and analyzed by WildFire, log data is sent back to the device after analysis, along with the analysis results. Logs are written to Monitor > Logs > WildFire.
- WildFire API—The WildFire license provides access to the WildFire API, which allows for programmatic access to the WildFire service in the cloud, apart from use by Palo Alto Networks firewalls. The WildFire API can be used to submit, analyze, and review reports for files submitted to the WildFire system. You can upload up to 100 files per day and query the WildFire portal up to 1000 times per day.

Results of the detailed analysis of the submitted files are also available through the WildFire portal and you can view this information without a subscription. You can use the WildFire portal to see which users were targeted, the applications that were used, and the malicious behavior that was observed. You can also configure the WildFire portal to send email notifications when results are available for review. Refer to "Using the WildFire Portal" on page 425.

Setting Up WildFire on the Firewall

Perform these tasks to set up your environment to use WildFire.

- 1. On the firewall, configure WildFire settings on the **Device > Setup** page. Refer to "Configuring WildFire Settings on the Firewall" on page 423.
- 2. On the firewall, configure your file blocking profiles to include the **forward** or **continueand-forward** action. Refer to "Configuring WildFire Forwarding" on page 424 for procedures, or refer to "File Blocking Profiles" on page 220 for more information.
- 3. Incorporate the file blocking profiles in a security policy, as you would for any other file blocking profiles. Refer to "Security Policies" on page 187.
- 4. Access the WildFire portal and configure optional settings. Refer to "Using the WildFire Portal" on page 425.

You can now access the WildFire portal to view reports. Refer to "Viewing WildFire Reports" on page 427.

Configuring WildFire Settings on the Firewall

► Device > Setup > WildFire

Use the WildFire tab to control the information to be sent to the WildFire server.



Note: To forward decrypted content to WildFire, you need to select the "Allow Forwarding of Decrypted Content" check box in **Device > Setup > Content-ID > URL Filtering** Settings box.

Table 158. WildFire Settings on the Firewall

Field	Description		
General Settings			
WildFire Server	Specify the URL of a WildFire server. Specify the value default-cloud to allow the firewall to automatically find the closest WildFire server.		
Maximum File Size (MB)	Specify the maximum file size that will be forwarded to the WildFire server (range 1-10 MB, default 2 MB). Files larger than the specified size are not sent.		

Field	Description		
Session Information Settings			
Settings	Specify the information to be forwarded to the WildFire server. By default, all are selected:		
	• Source IP—Source IP address that sent the suspected file.		
	• Source Port—Source port that sent the suspected file.		
	• Destination IP—Destination IP address for the suspected file.		
	• Destination Port—Destination port for the suspected file.		
	• Vsys—Firewall virtual system that identified the possible malware.		
	• Application—User application that was used to transmit the file.		
	• User—Targeted user.		
	• URL—URL associated with the suspected file.		
	• Filename—Name of the file that was sent.		

Table 158. WildFire Settings on the Firewall (Continued)

Configuring WildFire Forwarding

After configuring the WildFire settings on the firewall, you are now ready to configure forwarding to allow files to be sent to the WildFire system for analysis. You first create a file blocking profile and then include that profile in a security policy.

To configure a file blocking profile for WildFire:

- 1. Navigate to **Objects > Security Profiles > File Blocking**.
- 2. Click Add to add a new profile and enter a Name and Description.
- 3. Click **Add** in the **File Blocking Profile** window. Click in the **Names** field and enter a rule name.
- 4. Select the **Application(s)** that will be identified for this profile. For example, if you choose the application **web-browsing**, the profile will identify files downloaded when the user downloads files from a web page.
- 5. In the **File Type** field select the file types that you would like to forward for analysis.
- 6. In the **Direction** field, select **upload**, **download**, or **both**. If you select **both**, this will identify files that are uploaded or downloaded by the user.
- 7. In the **Action** field, select **forward**. This will cause any identified file to be forwarded to the WildFire system, analysis will occur and the file will then be delivered to the user. If you select **continue-and-forward**, the user will be prompted with a continue page before the download will occur.

If WildFire determines that the file is a malware file, a report will be generated on the WildFire Portal and a new antivirus signature will be created within 24-48 hours. If you are a WildFire subscriber, you will also see a **WildFire** log entry on the firewall and a signature will be created within an hour.

To configure a security policy for WildFire:

1. Navigate to **Policies** > **Security.**

- 2. Click on Add to create a new policy, or select an existing policy.
- 3. Click the **Actions** tab and under **Profile Setting** click the drop-down next to **File Blocking** and choose the Security Profile that you created.
- 4. Click **Commit** to active your changes.



Note: When you create a file blocking profile with the action **continue** or **continue-and-forward** (used for WildFire forwarding), you can only choose the application **web-browsing**. If you choose any other application, traffic that matches the security policy will not flow through the firewall due to the fact that the users will not be prompted with a continue page.

WildFire Data Filtering Log

This section describes the various forwarding actions that you will see in the logs when WildFire is configured and files are forwarded for analysis. The logs will be located in Monitor > Logs > Data Filtering.

Log	Description
wildfire-upload-success	The file was sent to the cloud. This means the file is not signed by a trusted file signer, it has not been previously analyzed by WildFire.
wildfire-upload-skip	This action will be displayed for all files identified as eligible to be sent to WildFire by a file blocking profile/security policy, but did not need to be analyzed by WildFire because it has already been analyzed previously. In this case, the "forward" action will appear in the Data Filtering log because it was a valid forward action, but it was not sent to WildFire and analyzed because the file has already been sent to the WildFire cloud from another session, possibly from another firewall.

Table 159. WildFire Data Filtering Log Descriptions

Using the WildFire Portal

To access the WildFire portal, go to *https://wildfire.paloaltonetworks.com* and log in using your Palo Alto Networks support credentials or your WildFire account.

The portal opens to display the dashboard, which lists summary report information for all of the firewalls associated with the specific WildFire license or support account (as well as any files that have been uploaded manually). The display includes the number of analyzed files and indicates how many are infected with malware, are benign, or are pending analysis.



Note: To upload a file manually, click *Upload File* in the upper right corner of the WildFire page.

				Doe	, John (Settings Log
Dashboard				Rep	orts Upload Fil
day					
Wildfire Stats	Device	Malware	Benign	Pending	Registered
	0003C103099	477	139	0	09/01/2011 08:27:37
	0004A100237	21	97	0	08/22/2011 11:06:37
	0006C103719	17	0	0	08/28/2011 01:30:11
	0006C104026	5	0	0	09/06/2011 06:18:2
	00060105153	0	3	0	09/08/2011 07:50:55
Malware Benign Pending	000002102102				
Malware Benign Pending	20000-100100				
Malware Benign Pending ays Wildfire Stats	Device	Malware	Benign	Pending	Registered
Malware Benign Pending ays Wildfire Stats	Device 0001A100029	Malware 9	Benign 21	Pending	Registered 09/07/2011 01:41:4
Malware Benign Pending ays Wildfire Stats	Device 0001A100029 0003C103099	Malware 9 3340	Benign 21 1074	Pending 0 0	Registered 09/07/2011 01:41:4 09/01/2011 08:27:3
ays Wildfire Stats	Device 0001A100029 0003C103099 0004A100237	Malware 9 3340 111	Benign 21 1074 240	Pending 0 0	Registered 09/07/2011 01:41:4 09/01/2011 08:27:3 08/22/2011 11:06:3
Malware Benign Pending ays Wildfire Stats	Device 0001A100029 0003C103098 0004A100237 0006C100450	Malware 9 3340 1111 0	Benign 21 1074 240 6	Pending 0 0 0 0	Registered 09/07/2011 01:41:4 09/01/2011 08:27:3 08/22/2011 11:06:3 09/11/2011 12:18:5
ays Wildfire Stats	Device 0001A100029 0004A100237 0006C100459 0006C100459	Malware 9 3340 111 0 0	Benign 21 1074 240 6 6	Pending 0 0 0 0 0	Registered 09/07/2011 01:41:4 09/01/2011 08:27:3 08/22/2011 11:06:3 09/11/2011 12:18:5 09/08/2011 10:51:2
ays	Device 0001A100029 0004A100237 0006C100450 0006C100453 0006C100588	Malware 9 3340 111 0 0 0	Benign 21 1074 240 6 6 1	Pending 0 0 0 0 0 0 0 0 0	Registered 09/07/2011 01:41:4 09/01/2011 08:27:3 08/22/2011 11:06:3 09/11/2011 12:18:50 09/08/2011 10:51:2 09/08/2011 10:51:2
ays	Device 0001A100029 0004A100237 0006C100450 0006C100453 0006C100588 0006C103719	Malware 9 3340 1111 0 0 0 19	Benign 21 1074 240 6 6 1 1 0	Pending 0 0 0 0 0 0 0 0 0 0	Registered 09/07/2011 01:41:4 09/07/2011 08:27:3 08/22/2011 11:06:3 09/11/2011 12:18:5 09/08/2011 10:51:2 09/08/2011 10:51:2 09/03/2011 08:53:1 08/28/2011 01:30:1
ays	Device 0001A100029 0004A100237 0006C100450 0006C100450 0006C100388 0006C103719 0006C103940	Malware 9 3340 111 0 0 0 19 4	Benign 21 1074 240 6 6 1 0 3	Pending 0 0 0 0 0 0 0 0 0 0 0 0 0	Registered 09/07/2011 01:41:4 09/07/2011 08:27:3 08/22/2011 11:06:33 09/11/2011 12:18:50 09/08/2011 10:51:2 09/03/2011 08:53:1 08/28/2011 01:30:1 09/07/2011 09:37:3
Malware Benign Pending Pays	Device 0001A100029 0006C100450 0006C100450 0006C100453 0006C100453 0006C100453 0006C100453 0006C103719 0006C10463	Malware 9 3340 111 0 0 0 19 4 23	Benign 21 1074 240 6 6 1 0 3 0	Pending 0 0 0 0 0 0 0 0 0 0 0 0	Registered 09/07/2011 01:41:42 09/07/2011 01:41:42 09/07/2011 08:27:33 09/11/2011 12:18:56 09/08/2011 10:51:22 09/08/2011 08:53:12 09/08/2011 01:30:11 09/07/2011 09:37:30 09/06/2011 06:18:22
Malware Benign Pending Pays	Device 00014100029 00056100450 00066100450 00066100463 00066100588 00066103719 00066103719 00066103940 00066104028	Malware 9 3340 111 0 0 0 19 4 23 0	Benign 21 1074 240 6 6 1 0 3 0 3 3	Pending 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	Registered 09/07/2011 01:41:42 09/07/2011 08:27:37 08/22/2011 11:06:37 09/11/2011 12:18:56 09/08/2011 10:51:22 09/08/2011 01:30:11 09/07/2011 09:37:34 09/06/2011 06:18:22 09/08/2011 07:50:55

Figure 55. WildFire Dashboard

Configuring Settings on the WildFire Portal

Click the **Settings** link at the top of the WildFire portal to configure the following settings:

Table 160. Settings on the WildFire Portal

Field	Description	
Password	To change your password, enter values in the following fields:	
	• Current Password—Enter your current password.	
	• New Password/Confirm Password—Enter and then reenter a new password.	
Time Zone	Select the time zone from the drop-down list. This is the time zone that is used to indicate when WildFire receives files.	

Field	Description
Email Notifications	Choose the email notifications that you would like to receive. The email notifications are sent to the currently logged in WildFire user. For each device, and for files that are manually uploaded to the WildFire server, you can choose any of the following email notifications:
	 Malware—Select this check box to have an Email notification sent when the file is determined to be malware.
	 Benign—Select this check box to have an Email notification sent when the file is determined to be benign.
	<i>Note:</i> You can also select both check boxes to receive Email notifications for all files (malware/benign).

Table 160. Settings on the WildFire Portal (Continued)

Viewing WildFire Reports

Click the **Reports** button at the top of the WildFire portal to view the list of available reports. Search options are available at the top of the page, and pagination controls are included. To view an individual report, click the 📧 icon to the left of the report name. To print a detailed report, use the print option on your browser.

	Reports			Dashb	oard Upload File
Sea	irch			Source Type 0004A100237	Search
_				Showing 1 - 50 of 68	6 first prev <u>next</u> <u>last</u>
	Received Time	Source	Filename	<u>Uri</u>	Verdict
P	09/12/2011 04:05 PM	0004A100237	HP_CLJ3600_32bit_HB.exe	unknown	Benign
P	09/12/2011 02:53 PM	0004A100237	DJ_SF_05_D2600_NonNet_Basic_Win	_WW_140_049.exe unknown	Benign
P	09/12/2011 01:57 PM	0004A100237	SetupEpicPlay.exe	d1.epicplay.com/aj/bundle/392	Benign
Þ	09/12/2011 12:51 PM	0004A100237	A11GX620.EXE	unknown	Malware
P	09/12/2011 12:46 PM	0004A100237	XvidSetup.exe	origin-ics.fivemillionfriends.com/IC/GPLAppBundler41/22	596/0/df Malware
Þ	09/12/2011 12:42 PM	0004A100237	vic.exe	us.f820.mail.yahoo.com/ya/upload_with_cred?cred=bQk1	WJkc1I95MT4 Benign
Þ	09/12/2011 12:26 PM	0004A100237	OJProL7X00_Basic_14.exe	unknown	Benign

Figure 56. WildFire Reports Page

Using the WildFire Portal

Appendix A Custom Pages

Custom response pages allow you to notify end users of policy violations and special access conditions. Each page can include references to the user's IP address, the URL for which access is attempted, and the URL category. These parameters can also be used in links to trouble-ticketing systems.

This appendix provides HTML code for the following default custom response pages:

- "Default Antivirus Response Page" in the next section
- "Default Application Block Page" on page 431
- "Default File Blocking Block Page" on page 431
- "Default URL Filtering Response Page" on page 432
- "Default Anti-spyware Download Response Page" on page 433
- "Default Decryption Opt-out Response Page" on page 433
- "Captive Portal Comfort Page" on page 434
- "URL Filtering Continue and Override Page" on page 434
- "SSL VPN Login Page" on page 435
- "SSL Certificate Revoked Notify Page" on page 436



Note: For information on importing and exporting custom response pages, refer to "Defining Custom Response Pages" on page 115.

Default Antivirus Response Page

<html>

```
<head>
<meta http-equiv=Content-Type content="text/html; charset=windows-1252">
<meta name=Generator content="Microsoft Word 11 (filtered)">
<title>This is a test</title>
<style>
<!--
/* Font Definitions */
@font-face
        {font-family:"Microsoft Sans Serif";</pre>
```

```
panose-1:2 11 6 4 2 2 2 2 2 4;}
 /* Style Definitions */
 p.MsoNormal, li.MsoNormal, div.MsoNormal
    {margin:0in;
    margin-bottom:.0001pt;
    font-size:12.0pt;
    font-family:"Times New Roman"; }
h4
    {margin-top:12.0pt;
    margin-right:0in;
    margin-bottom:3.0pt;
    margin-left:0in;
    page-break-after:avoid;
    font-size:14.0pt;
font-family:"Times New Roman";}
p.SanSerifName, li.SanSerifName, div.SanSerifName
    {margin:0in;
    margin-bottom:.0001pt;
    text-autospace:none;
    font-size:10.0pt;
    font-family: "Microsoft Sans Serif";
    font-weight:bold;}
p.BoldNormal, li.BoldNormal, div.BoldNormal
    {margin:0in;
    margin-bottom:.0001pt;
    font-size:12.0pt;
    font-family:"Times New Roman";
    font-weight:bold;}
span.Heading10
     {color:black
    font-weight:bold; }
p.SubHeading1, li.SubHeading1, div.SubHeading1
     {margin-top:12.0pt;
    margin-right:0in;
    margin-bottom:3.0pt;
    margin-left:0in;
    page-break-after:avoid;
    font-size:12.0pt;
    font-family:"Times New Roman";
    font-weight:bold;}
@page Section1
    {size:8.5in 11.0in;
    margin:1.0in 1.25in 1.0in 1.25in;}
div.Section1
    {page:Section1;}
- - >
</style>
</head>
<body lang=EN-US>
<div class=Section1>
This is a test.
</div>
</body>
</html>
```

Default Application Block Page

```
<html>
<head>
<title>Application Blocked</title>
<style>
#content{border:3px solid#aaa;background-
color:#fff;margin:40;padding:40;font-family:Tahoma,Helvetica,Arial,sans-
serif;font-size:12px;}
  h1{font-size:20px;font-weight:bold;color:#196390;}
  b{font-weight:bold;color:#196390;}
</style>
</head>
<body bgcolor="#e7e8e9"><div id="content">
<h1>Application Blocked</h1>
Access to the application you were trying to use has been blocked in
accordance with company policy. Please contact your system administrator if
you believe this is in error. 
<b>User:</b> <user/> 
<b>Application:</b> <appname/> 
</div>
</body>
</html>
```

Default File Blocking Block Page

<html>

```
<head>
<meta http-equiv=Content-Type content="text/html; charset=windows-1252">
<meta name=Generator content="Microsoft Word 11 (filtered)">
<title>This is a test</title>
<style>
<!--
 /* Font Definitions */
 @font-face
    {font-family:"Microsoft Sans Serif";
    panose-1:2 11 6 4 2 2 2 2 2 4;}
 /* Style Definitions */
 p.MsoNormal, li.MsoNormal, div.MsoNormal
    {margin:0in;
    margin-bottom:.0001pt;
    font-size:12.0pt;
    font-family:"Times New Roman";}
h4
    {margin-top:12.0pt;
    margin-right:0in;
    margin-bottom:3.0pt;
    margin-left:0in;
    page-break-after:avoid;
    font-size:14.0pt;
    font-family:"Times New Roman";}
p.SanSerifName, li.SanSerifName, div.SanSerifName
    {margin:0in;
    margin-bottom:.0001pt;
    text-autospace:none;
    font-size:10.0pt;
    font-family:"Microsoft Sans Serif";
    font-weight:bold;}
p.BoldNormal, li.BoldNormal, div.BoldNormal
    {margin:0in;
    margin-bottom:.0001pt;
    font-size:12.0pt;
    font-family: "Times New Roman";
    font-weight:bold;}
```

```
span.Heading10
    {color:black
    font-weight:bold;}
p.SubHeading1, li.SubHeading1, div.SubHeading1
    {margin-top:12.0pt;
    margin-right:0in;
    margin-bottom:3.0pt;
    margin-left:0in;
    page-break-after:avoid;
    font-size:12.0pt;
    font-family:"Times New Roman";
    font-weight:bold;}
@page Section1
    {size:8.5in 11.0in;
    margin:1.0in 1.25in 1.0in 1.25in;}
div.Section1
   {page:Section1;}
</style>
</head>
<body lang=EN-US>
<div class=Section1>
This is a test.
</div>
</body>
</html>
```

Default URL Filtering Response Page

```
<html>
<head>
<title>Web Page Blocked</title>
<style>
#content{border:3px solid#aaa;background-
color:#fff;margin:40;padding:40;font-family:Tahoma,Helvetica,Arial,sans-
serif;font-size:12px;}
 h1{font-size:20px;font-weight:bold;color:#196390;}
 b{font-weight:bold;color:#196390;}
</style>
</head>
<body bgcolor="#e7e8e9">
<div id="content">
<h1>Web Page Blocked</h1>
Access to the web page you were trying to visit has been blocked in
accordance with company policy. Please contact your system administrator if
you believe this is in error.
<b>User:</b> <user/> 
<b>URL:</b> <url/> 
<b>Category:</b> <category/> 
</div>
</body>
</html>
```
Default Anti-spyware Download Response Page

```
<application-type>
    <category>
         <entry name="networking" id="1">
              <subcategory>
                   <entry name="remote-access" id="1"/>
<entry name="proxy" id="2"/>
                   <entry name="encrypted-tunnel" id="3"/>
                   <entry name="routing" id="4"/>
                   <entry name="infrastructure" id="5"/>
<entry name="ip-protocol" id="6"/>
              </subcategory>
         </entry>
         <entry name="collaboration" id="2">
              <subcategory>
                   <entry name="email" id="7"/>
                   <entry name="instant-messaging" id="8"/>
                   <entry name="social-networking" id="9"/>
                   <entry name="internet-conferencing" id="10"/>
                   <entry name="voip-video" id="11"/>
              </subcategory>
         </entry>
         <entry name="media" id="3">
              <subcategory>
                   <entry name="video" id="12"/>
                   <entry name="gaming" id="13"/>
                   <entry name="audio-streaming" id="14"/>
              </subcategory>
         </entry>
         <entry name="business-systems" id="4">
              <subcategory>
                   <entry name="auth-service" id="15"/>
                   <entry name="database"id="16"/>
                   <entry name="erp-crm" id="17"/>
                   <entry name="general-business" id="18"/>
                   <entry name="management" id="19"/>
                   <entry name="office-programs" id="20"/>
<entry name="software-update" id="21"/>
                   <entry name="storage-backup" id="22"/>
              </subcategory>
          </entry>
          <entry name="general-internet" id="5">
              <subcategory>
                   <entry name="file-sharing" id="23"/>
                   <entry name="internet-utility" id="24"/>
              </subcategory>
         </entry>
    </category>
    <technology>
           <entry name="network-protocol" id="1"/>
           <entry name="client-server" id="2"/>
           <entry name="peer-to-peer" id="3"/>
           <entry name="web-browser" id="4"/>
    </technology>
</application-type>
```

Default Decryption Opt-out Response Page

```
<h1>SSL Inspection</h1>
In accordance with company security policy, the SSL encrypted connection
you have initiated will be temporarily unencrypted so that it can be
inspected for viruses, spyware, and other malware.
After the connection is inspected it will be re-encrypted and sent to its
destination. No data will be stored or made available for other purposes.
<b>IP:</b> <url> 
<b>Category:</b> <category/>
```

Captive Portal Comfort Page

```
<h1 ALIGN=CENTER>Captive Portal</h1>
<h2 ALIGN=LEFT>In accordance with company security policy, you have to authenticate before accessing the network.</h2>
<pan form/>
```

URL Filtering Continue and Override Page

```
<html>
<head>
<title>Web Page Blocked</title>
<style>
#content{border:3px solid#aaa;background-
color:#fff;margin:40;padding:40;font-family:Tahoma,Helvetica,Arial,sans-
serif;font-size:12px;}
  h1{font-size:20px;font-weight:bold;color:#196390;}
  b{font-weight:bold;color:#196390;}
       form td, form input {
               font-size: 11px;
               font-weight: bold;
       #formtable {
               height: 100%;
               width: 100%;
       #formtd {
               vertical-align: middle;
       #formdiv {
               margin-left: auto;
               margin-right: auto;
       }
</style>
<script type="text/javascript">
function pwdCheck() {
    if(document.getElementById("pwd")) {
        document.getElementById("continueText").innerHTML = "If you require
access to this page, have an administrator enter the override password
here:";
    }
}
</script>
</head>
<body bgcolor="#e7e8e9">
<div id="content">
<h1>Web Page Blocked</h1>
Access to the web page you were trying to visit has been blocked in
accordance with company policy. Please contact your system administrator if
you believe this is in error.
<b>User:</b> <user/> 
<b>URL:</b> <url/> 
<b>Category:</b> <category/> 
<hr>>
If you feel this page has been incorrectly blocked, you
may click Continue to proceed to the page. However, this action will be
logged.
<div id="formdiv">
<pan_form/>
</div>
<a href="#" onclick="history.back();return false;">Return to previous page
a>
</div>
</body>
</html>
```

SSL VPN Login Page

```
<HTML>
<HEAD>
<TITLE>Palo Alto Networks - SSL VPN</TITLE>
<meta http-equiv="Content-Type" content="text/html; charset=iso-8859-1">
<link rel="stylesheet" type="text/css" href="/styles/</pre>
falcon content.css?v=@@version">
<style>
td {
   font-family: Verdana, Arial, Helvetica, sans-serif;
font-weight: bold;
   color: black; /*#FFFFFF; */
}
.msg {
   background-color: #ffff99;
   border-width: 2px;
   border-color: #ff0000;
   border-style: solid;
   padding-left: 20px;
   padding-right: 20px;
   max-height: 150px;
   height: expression( this.scrollHeight > 150 ? "150px" : "auto" ); /* sets
max-height for IE */
   overflow: auto;
}
.alert {font-weight: bold;color: red;}
</style>
</HEAD>
<BODY bgcolor="#F2F6FA">
    bottom: 2px solid #888888;">
       background-repeat: no-repeat">
            
        <div align="center">
       <hi>Palo Alto Networks - SSL VPN Portal</hi>
    </div>
<div id="formdiv">
<pan form/>
</div>
</BODY>
</HTML>
```

SSL Certificate Revoked Notify Page

```
<META HTTP-EQUIV="Content-Type" CONTENT="text/html; charset=iso-8859-1">
<html>
<head>
<title>Certificate Error</title>
<style>
#content{border:3px solid#aaa;background-
color:#fff;margin:40;padding:40;font-family:Tahoma,Helvetica,Arial,sans-
serif;font-size:12px;}
 h1{font-size:20px;font-weight:bold;color:#196390;}
 b{font-weight:bold;color:#196390;}
</style>
</head>
<body bgcolor="#e7e8e9">
<div id="content">
<h1>Certificate Error</h1>
There is an issue with the SSL certificate of the server you are trying to
contact.
<b>Certificate Name:</b> <certname/> 
<b>IP:</b> <url/> 
<b>Issuer:</b> <issuer/> 
<b>Status:</b> <status/> 
<b>Reason:</b> <reason/> 
</div>
</body>
</html>
```

Appendix B Application Categories, Subcategories, Technologies, and Characteristics

The appendix lists application-related categories defined by Palo Alto Networks:

- "Application Categories and Subcategories" in the next section
- "Application Technologies" on page 439
- "Application Characteristics" on page 439

The Applipedia database can also be found on the firewall in **Objects > Applications**, online at http://apps.paloaltonetworks.com/applipedia/ and an App is available on the Apple App Store.

Application Categories and Subcategories

The following application categories and subcategories are supported:

- business-system
 - auth-service
 - database
 - erp-crm
 - general-business
 - management
 - office-programs
 - software-update
 - storage-backup
- collaboration
 - email
 - instant-messaging
 - internet-conferencing

- social-business
- internet-utility
- social-networking
- voip-video
- web-posting
- general-internet
 - file-sharing
 - internet-utility
- media
 - audio-streaming
 - gaming
 - photo-video
- networking
 - encrypted-tunnel
 - infrastructure
 - ip-protocol
 - proxy
 - remote-access
 - routing
- unknown

Application Technologies

The following application technologies are supported.

Table 161. Application Technologies

ltem	Description
browser-based	An application that relies on a web browser to function.
client-server	An application that uses a client-server model where one or more clients communicate with a server in the network.
network-protocol	An application that is generally used for system to system communication that facilitates network operation. This includes most of the IP protocols.
peer-to-peer	An application that communicates directly with other clients to transfer information instead of relying on a central server to facilitate the communication.

Application Characteristics

The following application characteristics are supported.

Table 162. Application Characteris	stics
------------------------------------	-------

ltem	Description
Evasive	Uses a port or protocol for something other than its originally intended purpose with the hope that it will traverse a firewall.
Excessive Bandwidth	Consumes at least 1 Mbps on a regular basis through normal use.
Prone to Misuse	Often used for nefarious purposes or is easily set up to expose more than the user intended.
Transfers Files	Has the capability to transfer a file from one system to another over a network.
Tunnels Other Apps	Is able to transport other applications inside its protocol.
Used by Malware	Malware has been known to use the application for propagation, attack, or data theft, or is distributed with malware.
Vulnerability	Has publicly reported vulnerabilities.
Widely Used	Likely has more than 1,000,000 users.
Continue Scanning for Other Applications	Instructs the firewall to continue looking to see if other application signatures match. If this option is not selected, the first matching signature is reported and the firewall stops looking for additional matching applications.

Application Characteristics

Appendix C Common Criteria/Federal Information Processing Standards Support

You can configure the firewall to support the Common Criteria Evaluation Assurance Level 4+ (CCEAL4+) and the Federal Information Processing Standards 140-2 (FIPS 140-2), which are security certifications that ensure a standard set of security assurances and functionalities. These certifications are often required by civilian U.S. government agencies and government contractors.

For more details on these certifications, refer to the following documents:

- FIPS—<u>http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140sp/</u> 140sp1877.pdf
- CCEAL4+—<u>http://www.niap-ccevs.org/st/vid10392/</u>

Note: PAN-OS 5.0 certification is in process. Visit http://csrc.nist.gov/groups/STM/cmvp/ documents/140-1/140InProcess.pdf and search for Palo Alto Networks in the Cryptographic Module Validation Program FIPS 140-1 and FIPS 140-2 Modules In Process List for details.

Enabling CC/FIPS Mode

Use the following procedure to enable CC/FIPS mode on a software version that supports CC/FIPS. Keep in mind that when you enable CC/FIPS, the device will be reset the factory default settings; all configuration will be removed.

- 1. Boot the firewall into maintenance mode as follows:
 - a. Establish a serial connection with the firewall.
 - b. Reboot the device and press the **m** key on the keyboard when you see the following prompt: **Press m to boot to maint partition**.
 - c. Press any key on your keyboard when prompted to stop the automatic boot, and then select **PANOS (maint)** as the booting partition.

For more detailed information about maintenance mode, refer to the *PAN-OS Command Line Interface Reference Guide*.

- 2. Select Set CCEAL4 Mode from the menu.
- 3. Select Enable CCEAL4 Mode from the menu.

4. When prompted, select **Reboot**.

After successfully switching to CC/FIPS mode, the following status displays: **CCEAL4 mode enabled successfully**. In addition, **CC** will display at all times in the status bar at the bottom of the web interface. In addition, the console port will now be available as a status output port only. In addition, the default admin login credentials change to admin/ paloalto.

CC/FIPS Security Functions

When CC/FIPS is enabled, the following apply:

- To log into the firewall, the browser must be TLS 1.0 compatible.
- All passwords on the firewall must be at least six characters.
- Accounts are locked after the number of failed attempts that is configured on the **Device > Setup > Management** page. If the firewall is not in CC/FIPS mode, it can be configured so that it never locks out; however in CC/FIPS mode, and lockout time is required.
- The firewall automatically determines the appropriate level of self-testing and enforces the appropriate level of strength in encryption algorithms and cipher suites.
- Non-CC/FIPS approved algorithms are not decrypted and are thus ignored during decryption.
- When configuring IPSec, a subset of the normally available cipher suites is available.
- Self-generated and imported certificates must contain public keys that are 2048 bits (or more).
- The serial port is disabled.
- Telnet, TFTP, and HTTP management connections are unavailable.
- Surf control is not supported.
- High availability (HA) encryption is required.
- PAP authentication is disabled.

Appendix D Open Source Licenses

The software included in this product contains copyrighted software that is licensed under the General Public License (GPL). A copy of that license is included in this document. You may obtain the complete Corresponding Source code from us for a period of three years after our last shipment of this product by sending a money order or check for \$5 to:

Palo Alto Networks Open Source Request 4401 Great America Parkway Santa Clara, Ca. 95054

Some components of this product may be covered under one or more of the open source licenses listed in this appendix:

- "Artistic License" on page 446
- "BSD" on page 447
- "GNU General Public License" on page 448
- "GNU Lesser General Public License" on page 452
- "MIT/X11" on page 458
- "OpenSSH" on page 458
- "PSF" on page 462
- "PHP" on page 462
- "Zlib" on page 463

Artistic License

This document is freely plagiarized from the 'Artistic License', distributed as part of the Perl v4.0 kit by Larry Wall, which is available from most major archive sites

This documents purpose is to state the conditions under which these Packages (See definition below) viz: "Crack", the Unix Password Cracker, and "CrackLib", the Unix Password Checking library, which are held in copyright by Alec David Edward Muffett, may be copied, such that the copyright holder maintains some semblance of artistic control over the development of the packages, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications. Definitions:

A "Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification, or segments thereof.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.

2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.

3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when AND WHY you changed that file, and provided that you do at least ONE of the following:

a) place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as uunet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.

b) use the modified Package only within your corporation or organization.

c) rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide separate documentation for each non-standard executable that clearly documents how it differs from the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:

a) distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.

b) accompany the distribution with the machine-readable source of the Package with your modifications.

c) accompany any non-standard executables with their corresponding Standard Version executables, giving the non-standard executables non-standard names, and clearly documenting the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. YOU MAY NOT CHARGE A FEE FOR THIS PACKAGE ITSELF. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that YOU DO NOT ADVERTISE this package as a product of your own.

6. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

7. THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.

BSD

The following copyright holders provide software under the BSD license:

- Julian Steward
- Thai Open Source Software Center Ltd
- The Regents of the University of California
- Nick Mathewson
- Niels Provos
- Dug Song
- Todd C. Miller
- University of Cambridge
- Sony Computer Science Laboratories Inc.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

GNU General Public License

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble:

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/ OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

GNU Lesser General Public License

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble:

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

* a) The modified work must itself be a software library.

* b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.

* c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.

* d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

* a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

* b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

* c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

* d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

* e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

* a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

* b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license

would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

MIT/X11

Copyright (C) 2001-2002 Daniel Veillard. All Rights Reserved.

Copyright (C) 2001-2002 Thomas Broyer, Charlie Bozeman and Daniel Veillard. All Rights Reserved.

Copyright (C) 1998 Bjorn Reese and Daniel Stenberg.

Copyright (C) 2000 Gary Pennington and Daniel Veillard.

Copyright (C) 2001 Bjorn Reese <breese@users.sourceforge.net>

Copyright (c) 2001, 2002, 2003 Python Software Foundation

Copyright (c) 2004-2008 Paramjit Oberoi cparam.cs.wisc.edu>

Copyright (c) 2007 Tim Lauridsen <tla@rasmil.dk>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

OpenSSH

This file is part of the OpenSSH software.

The licences which components of this software fall under are as follows. First, we will summarize and say that all components are under a BSD licence, or a licence more free than that.

OpenSSH contains no GPL code.

1) Copyright (c) 1995 Tatu Ylonen <ylo@cs.hut.fi>, Espoo, Finland

All rights reserved

As far as I am concerned, the code I have written for this software can be used freely for any purpose. Any derived versions of this software must be clearly marked as such, and if the derived work is incompatible with the protocol description in the RFC file, it must be called by a name other than "ssh" or "Secure Shell".

[Tatu continues]

However, I am not implying to give any licenses to any patents or copyrights held by third parties, and the software includes parts that are not under my direct control. As far as I know, all included source code is used in accordance with the relevant license agreements and can be used freely for any purpose (the GNU license being the most restrictive); see below for details.

[However, none of that term is relevant at this point in time. All of these restrictively licenced software components which he talks about have been removed from OpenSSH, i.e.,

-RSA is no longer included, found in the OpenSSL library

-IDEA is no longer included, its use is deprecated

-DES is now external, in the OpenSSL library

-GMP is no longer used, and instead we call BN code from OpenSSL

-Zlib is now external, in a library

-The make-ssh-known-hosts script is no longer included

-TSS has been removed

-MD5 is now external, in the OpenSSL library

-RC4 support has been replaced with ARC4 support from OpenSSL

-Blowfish is now external, in the OpenSSL library

[The licence continues]

Note that any information and cryptographic algorithms used in this software are publicly available on the Internet and at any major bookstore, scientific library, and patent office worldwide. More information can be found e.g. at "http://www.cs.hut.fi/crypto".

The legal status of this program is some combination of all these permissions and restrictions. Use only at your own responsibility. You will be responsible for any legal consequences yourself; I am not making any claims whether possessing or using this is legal or not in your country, and I am not taking any responsibility on your behalf.

NO WARRANTY

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,

REPAIR OR CORRECTION.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/ OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

2) The 32-bit CRC compensation attack detector in deattack.c was contributed by CORE SDI S.A. under a BSD-style license.

Cryptographic attack detector for ssh - source code

Copyright (c) 1998 CORE SDI S.A., Buenos Aires, Argentina.

All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that this copyright notice is retained.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES ARE DISCLAIMED. IN NO EVENT SHALL CORE SDI S.A. BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES RESULTING FROM THE USE OR MISUSE OF THIS SOFTWARE.

Ariel Futoransky <futo@core-sdi.com> <http://www.core-sdi.com>

3) ssh-keyscan was contributed by David Mazieres under a BSD-style license.

Copyright 1995, 1996 by David Mazieres <dm@lcs.mit.edu>.

Modification and redistribution in source and binary forms is permitted provided that due credit is given to the author and the OpenBSD project by leaving this copyright notice intact.

4) The Rijndael implementation by Vincent Rijmen, Antoon Bosselaers and Paulo Barreto is in the public domain and distributed with the following license:

@version 3.0 (December 2000)

Optimised ANSI C code for the Rijndael cipher (now AES)

@author Vincent Rijmen <vincent.rijmen@esat.kuleuven.ac.be>

@author Antoon Bosselaers <antoon.bosselaers@esat.kuleuven.ac.be>

@author Paulo Barreto <paulo.barreto@terra.com.br>

This code is hereby placed in the public domain.

THIS SOFTWARE IS PROVIDED BY THE AUTHORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

5) One component of the ssh source code is under a 3-clause BSD license, held by the University of California, since we pulled these parts from original Berkeley code.

Copyright (c) 1983, 1990, 1992, 1993, 1995

The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS

BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

6) Remaining components of the software are provided under a standard 2-term BSD licence with the following names as copyright holders:

-Markus Friedl

-Theo de Raadt

-Niels Provos

-Dug Song

-Aaron Campbell

-Damien Miller

-Kevin Steves

-Daniel Kouril

-Wesley Griffin

-Per Allansson

-Nils Nordman

-Simon Wilkinson

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. 1. This LICENSE AGREEMENT is between the Python Software Foundation ("PSF"), and the Individual or Organization ("Licensee") accessing and otherwise using Python 2.3 software in source or binary form and its associated documentation.

2. Subject to the terms and conditions of this License Agreement, PSF hereby grants Licensee a nonexclusive, royalty-free, world-wide license to reproduce, analyze, test, perform and/or display publicly, prepare derivative works, distribute, and otherwise use Python 2.3 alone or in any derivative version, provided, however, that PSF's License Agreement and PSF's notice of copyright, i.e., "Copyright (c) 2001, 2002, 2003 Python Software Foundation; All Rights Reserved" are retained in Python 2.3 alone or in any derivative version prepared by Licensee.

3. In the event Licensee prepares a derivative work that is based on or incorporates Python 2.3 or any part thereof, and wants to make the derivative work available to others as provided herein, then Licensee hereby agrees to include in any such work a brief summary of the changes made to Python 2.3.

4. PSF is making Python 2.3 available to Licensee on an "AS IS" basis. PSF MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, PSF MAKES NO AND DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF PYTHON 2.3 WILL NOT INFRINGE ANY THIRD PARTY RIGHTS. 5. PSF SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON 2.3 FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS A RESULT OF MODIFYING, DISTRIBUTING, OR OTHERWISE USING PYTHON 2.3, OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.

6. This License Agreement will automatically terminate upon a material breach of its terms and conditions.

7. Nothing in this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between PSF and Licensee. This License Agreement does not grant permission to use PSF trademarks or trade name in a trademark sense to endorse or promote products or services of Licensee, or any third party.

8. By copying, installing or otherwise using Python 2.3, Licensee agrees to be bound by the terms and conditions of this License Agreement.

PHP

The PHP License, version 3.01

Copyright (c) 1999 - 2009 The PHP Group. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. The name "PHP" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact group@php.net.

4. Products derived from this software may not be called "PHP", nor may "PHP" appear in their name, without prior written permission from group@php.net. You may indicate that your software works in conjunction with PHP by saying "Foo for PHP" instead of calling it "PHP Foo" or "phpfoo"

5. The PHP Group may publish revised and/or new versions of the license from time to time. Each version will be given a distinguishing version number. Once covered code has been published under a particular version of the license, you may always continue to use it under the terms of that version. You may also choose to use such covered code under the terms of any subsequent version of the license published by the PHP Group. No one other than the PHP Group has the right to modify the terms applicable to covered code created under this License.

6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes PHP software, freely available from <http://www.php.net/software/>". THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE PHP DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project, please see http://www.php.net>.

PHP includes the Zend Engine, freely available at <http://www.zend.com>.

Zlib

Copyright (C) 1995-2005 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.

2.Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.

3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org

Mark Adler madler@alumni.caltech.edu

Index

Α

access domains firewall 56, 61 Panorama 392 accounts authentication profiles 62 creating administrator 59 username and password requirements 57 acknowledging alarms 85 active configuration, updating 37, 41 Active Directory configuring User-ID Agent 298 installing User-ID Agent 297 uninstalling and upgrading User-ID Agent 301 User-ID Agent 296 active/active high availability 93 active/passive high availability 93 adding devices to Panorama 382 address groups, defining 230 addresses defining 227 defining address groups 230 defining group 230 defining ranges 227 administrator accounts, about 56 accounts, creating 59 authentication options 56 page lockout 61, 384, 391 profiles, about 56 roles, about 56 roles, defining 58 agent configuring terminal server 301 GlobalProtect 335 setting up GlobalProtect 353 User-ID 284 using GlobalProtect 353 aggregate Ethernet interfaces configuring 141 groups 141 aggregate groups 141 alarms acknowledged 85

alarm icon 85 log 68 log settings 73 making the icon visible 85 thresholds 178 unacknowledged 85 viewing 85 allow list URL filtering profile 219 wildcard patterns 218 anti-spyware profiles about 213 defining 213 antivirus profiles defining 212 settings 212 antivirus response pages 429 App-ID, requesting 277 Application Command Center (ACC), using 253 application exception policies 246 application exceptions 213 application groups, defining 238 application override policies about 205 defining 205 applications ACC page 254 categories 233, 437 characteristics 234, 439 custom with application override 205 defining 233 defining filters 238 defining groups 238 details 232 exceptions 213 filters 231 identifying unknown 276 response page 431 searching 232 sub category 233 subcategories 437 technologies 439 updating threat definitions 55 applications list 254 App-Scope reports

change monitor report 258 network monitor report 261 summary report 257 threat map report 259, 260, 263 viewing 256 **ARP** entries on L3 subinterfaces 144 on main L3 interfaces 138 on VLAN interfaces 134, 146 AS BGP 154 description 154 audit configuration 49 authentication GlobalProtect 336 **IKE 312** LDAP 56 local database 56 options for administrator 56 RADIUS 56 remote 30, 49 sequence 67 Authentication Header (AH) 312 authentication profiles about 62 Kerberos settings 67 LDAP settings 66 RADIUS settings 65 setting up 62 authentication sequences about 68 setting up 68

B

backing up firewall configurations 417 BGP about 153 configuring virtual router 163 description 154 redistribution profiles 158 virtual routers 160 block list URL filtering profile 218 wildcard patterns 218 blocking, file profiles 220 botnet infected clients 267 botnets about 267 reports 267 BrightCloud service 217 browsers, supported 27

С

candidate configuration about 37, 41 saving and rolling back 37, 41 captive portal 62 comfort page 115, 434

configuring firewall for 290 defining policies 206 settings 207, 209 certificate authority (CA) CRL 46 GlobalProtect 336 OCSP 46 trusted CA certificate 86 certificates CRL 46 decryption 86 exporting 87 generating 87 GlobalProtect 336 importing 87 OCSP 46 Panorama server 87 renew 87 revoke 86 trusted CA 86 web 86 clear text traffic, and QoS 356 clients botnet infected 267 downloading and activating GlobalProtect 352 clock, setting 30, 49 committing changes 25 options 25 Panorama 398 comparison of configurations 49 configuration audit 49 configuration bundle exports 417 configuration log about 68 defining remote logging 72, 73, 74 viewing 267 configuration management 37, 41 configuration, sample VPN 320 content-id settings 43 conventions, typographical 15 CPU utilization 252 CRL 46 crossover cables 100 crypto profiles 318, 319 crypto profiles, about 312 custom group reports 272 custom reports 274 custom signatures about 244 spyware 244 vulnerability 244

D

dashboard firewall 252 data filtering

ACC page 255 data patterns 243 defining profiles 223 HIP matches on ACC page 255 list 255 pattern settings 224 profile settings 223, 225 profiles 223 profiles and patterns 224 viewing logs 69, 266 data patterns adding new 240 data filtering profiles 224 defining 243 rules 240 data protection adding 45 changing password 45 dead peer protection 315 decoders and actions 212 decryption policy 248 defining configuration templates 419 denial of service (Dos), profiles 225 deployment options Layer 2 124 Layer 3 124 PPPoE 124 tap mode 125 virtual wire 120 deployment types 120 deployment, viewing information 416 deployments 120 device groups adding 386 object assignment and sharing 396 selecting 396 using 396 device priority, HA 378 devices adding 385 management 29 master 387 DHCP firewall options 171 relay 171 servers 171 settings 171, 356 Diffie-Hellman (DH) group 312, 318 discard options, DOS profiles 180 disk utilization 252 DNS servers 172 DNS proxy about 173 settings 173 do not fragment (DF) 181 domain name 30 DoS defining policies 209

profiles 225 protection profiles 225 downgrading software 53 duplex settings 128, 139, 149, 150 Duplicate Address Detection (DAD) 132, 136, 144 Dynamic Block Lists 242 dynamic updates about 55 scheduling 56 dynamic URL timeout 44

E

editing settings on a page 24 email scheduling report delivery 273 email notification settings defining 83, 85 in logging profiles 247 Encapsulating Security Payload (ESP) 312 encrypting private keys and passwords 90 Ethernet interfaces, configuring 142 exchange mode 314 exports certificates 86 configuration bundle 417 scheduling log 71

F

fail over 177 features and benefits 18 file blocking defining profiles 220 profiles, defining 246 settings 220 file blocking page 431 filters application 231, 238 sub category 231 FIPS 441 firewall configuring WildFire settings 423 features and benefits 18 introduction 17 latitude and longitude 31 navigating the user interface 26 User-ID Agent 286 using the web interface 23 flood protection 175 flood, zone protection settings 178, 358 FTP server, saving logs to 71

G

gateway about 335 GlobalProtect 335 setting up GlobalProtect 347 getting help 24

GlobalProtect about 335 agents 335 authentication 336 connection process 336 downloading and activating clients 352 gateways 335 large scale VPN 323 portal backup 334 portals 335 response page 115 setting up 337 setting up agents 353 setting up gateways 347 setting up portals 341 using the agent 353 groups aggregate interface 141 defining service 240 device 386

Η

HA interfaces 150 HA upgrading 51 HA1 and HA2 ports 101 hello interval, HA 378 help 24 high availability about 93 active/active 93 active/passive 93 configuration notes 108 configuring 101 configuring on Panorama 377 defining interfaces 150 enabling 101 interfaces 150 notes about setting up 100 Panorama 377 rules for operation and failover 93 setting up 99 upgrading PAN-OS software 51 hold time 378 Host Information Profile (HIP) HIP match log settings 73 match log 267 matches on ACC page 255 setting up 340 setting up objects 337 host name, defining 30, 49 HTML block pages 429

I

ICMP flood 178, 179 identification, IKE 312 IKE about 310, 312 AH 312

authentication 312 crypto profile settings 318 crypto profiles, about 312 dead peer protection 315 defining crypto profiles 318 DH group 312 ESP 312 exchange mode 314 identification 312 IKE gateways about 312 setting up 175, 314 settings 314 installation, Panorama 372 interface management profiles 175 interfaces aggregate groups 141 configuring aggregate Ethernet 141, 142 high availability 150 L2 subinterfaces 129, 149 L2, main 128 L3 subinterfaces 134 L3, main 130 summary of 127 tap 149 viewing status 128, 252 IPSec AH 312 crypto profile settings 319 defining crypto profiles 319 DH group 312 ESP 312 large scale VPN overview 323 lifetime 312 number of tunnels 311 setting up tunnels 315 IPv6 176 IPv6 addresses 228

Κ

Kerberos administrator roles 56 configuring server settings 67 knowledge base 117

L

L2 deployments, about 124 L2 interfaces configuring 128 configuring subinterfaces 129 main 128 subinterfaces 129, 149 L3 deployments, about 124 L3 interfaces configuring 130 configuring subinterfaces 134 loopback 146, 147 main 130
shared gateways 115 subinterfaces 134 Large Scale VPN overview 323 latitude and longitude 31 Layer 130 LDAP authentication 56 configuring server settings 66 licenses installing 50, 86 open source 445 link groups, HA 106 link speed and duplex 128, 139, 149, 150 link state setting 128, 139, 149, 150 viewing 252 local identification 314 lockout on Administrator's page 61, 384, 391 log destinations about 70 email 83, 85 SNMP traps 75 syslog 76 log exports 71 log forwarding defining profiles 247 profile settings 247 log page links 254 logs 266 about 68 alarms 68,73 clearing 74 configuration 68 configuration settings 72 defining remote logging for the configuration 72, 73, 74 for threat and traffic logs 246 HIP match 267 HIP match settings 73 links from ACC pages 254 managing 74 resolve hostname 265 saving to FTP server 71 scheduling exports 71 system 69 threat 69 viewing 264 viewing URL filtering 266 WildFire 70 loopback interfaces defining 146, 147 management port 30 LSA 153

Μ

management interface CLI 19

configuring 30, 49 options 19 Panorama 19 web 19 managing configurations 37, 41 master device 387 Master Key and Diagnostics page 90 MD5 162 memory utilization 252 MIBs 47,75 modifying settings on a page 24 monitor profiles 177 multicast routing about 154 settings 169 multiple virtual systems 31, 111, 113

Ν

NAT defining policies 194 NAT64 196 policies 190 policy examples 195 types 192 navigation 26 Netflow about 85 configuring 85 network profiles 174 network settings 30, 49 networking overview 120 next hop 157 NFS 373 external log storage 373 Panorama high availability 378 storage partitions 373 NIS servers 172 NSSA (not so stub area) 161 NT LAN Manager (NTLM) 208, 285 NTP servers 172

0

objects overview 226 sharing in Panorama 395 OCSP 46 open source licenses 445 Open Virtual Machine Format (OVF) 372 OSPF about 153 configuring virtual router 160 redistribution profiles 158 virtual routers 160

Ρ

packet capture 266 accessing 266

APP-ID 278 capture files 278 configuring capture settings 278 profile setting 212, 214, 215, 216 taking captures 278 packet-based attack protection 175, 180 Panorama ACC tab 382 access domains 392 accessing 382 adding devices 385 administrator account creation 389 administrator options 387 administrator roles 388 committing 398 configuration bundle exports 417 configuring IP address 32 dashboard 382 device tab 383 enabling access 32 hardware appliance 376 high availability 377 installing 372 interface description 382 logging 403 monitor tab 382 network tab 382 objects 395 objects tab 382 panorama tab 383 policies tab 382 server certificate 87 set up 371 shared policies, defining 382, 393 tab 383 tabs 382 templates 399 templates, configuring 401 upgrading software 417, 419 user account lockout 61, 384, 391 user interface 382 virtual appliance 372 PAN-OS software upgrading 50, 61, 392 version 252 PAN-OS, upgrading software 50 passive hold time, HA 378 passive link state 106 passive/active high availability 93 password data protection 45 encrypting 90 minimum password complexity 35 new 22 profiles 59 path groups, HA 379 PDF summary reports creating 271, 273 designing 271

displaying 271 viewing 270 peer identification 314 Perfect Forward Security (PFS) 312 policies about 183 about NAT 190 about policy based forwarding 199 about security 187 and virtual systems 111, 113 data patterns 243 defining address objects 227 defining captive portal 206 defining decryption 202 defining DoS 209 defining NAT 194 guidelines on defining 184 other policy objects 226 post rules 394 pre rules 393 QoS 359 shared 382, 393 specifying users and applications 186 types 183 virtual systems 110 policies shared global 393 policy based forwarding (PBF) about 199 and monitor profiles 177 defining 199 portal about 335 GlobalProtect 335 setting up GlobalProtect 341 post rules in policies 394 **PPPoE** about 124 deployments 124 settings 131 pre rules in policies 393 private key, encrypting 90 profile groups, defining 246 profiles about monitor 177 about security 211 anti-spyware 213 antivirus 212 antivirus, application exceptions 213 antivirus, decoders and actions 212 data filtering 223 defining log forwarding 247 file blocking 220, 246 flood protection 175 IKE crypto 318 IKE crypto profile settings 318 interface management 175 IPSec crypto 319 IPSec crypto profile settings 319 logging 246

network 174 packet-based attack 175 QoS 355, 358 reconnaissance detection 175 redistribution 158 security groups 211, 246 tunnel monitor 177 URL filtering 217 vulnerability protection 215, 217 zone protection 178, 358 proxy DNS, about 173

Q

QoS classes 358, 359 clear text traffic 356 configuration instructions 355 egress settings 358 marking 190 policies 359 priority settings 358 profiles 355, 358 settings 190 settings for firewall interfaces 355 tunneled traffic 356 quality of service (QoS) 355

R

RADIUS authentication 56 authentication profiles 62 defining server settings 65 random early drop 178 rebooting the device 30, 39, 49 rebooting the firewall 39 reconnaissance detection 175 redistribution profiles about 158 configuring 158 regions about 230 policies 230 regular expressions, data patterns 240 rematching sessions 45 remote authentication 30, 49 rendezvous point 169 reports App-Scope 256 creating custom group 272 custom 274 PDF summary 270 scheduling email delivery 273 top 50 273 user activity 272, 404 viewing 273 viewing WildFire 427 reports and logs

custom reports 274 identifying unknown applications 276 using the Application Command Center 253 using the dashboard 252 viewing App-Scope reports 256 viewing PDF summary reports 270 viewing reports 273, 274 Representational State Transfer (REST) 19 requesting support 117 required fields 26 resolve hostname 265 response pages antivirus 115, 429 application block 115, 431 captive portal 115, 434 defining 115 file blocking 115, 431 file blocking continue 115 GlobalProtect portal help 115 GlobalProtect portal login 115 SSL certificate errors notify page 116 SSL certificate revoked notify 436 SSL decryption opt-out 116 types 89, 115 URL filtering continue and override 116, 434 response thresholds 178, 179 restarting the firewall 39 RIP about 153 configuring virtual router 159 redistribution profiles 158 virtual routers 159 roles about 56 defining administrator 58 rolling back a candidate configuration 37, 41 Router Advertisement 133, 137, 145 routing about virtual routers and routing protocols 153 multicast 154 routing protocols about 153 BGP 160 **OSPF** 160 **RIP** 159 virtual routers 153 rules application exception policy 246 security policy 187

S

sample VPN configuration 320 saving a candidate configuration 37, 41 schedules configuration bundle exports 417 defining 246, 250 security

defining profile groups 211, 246 profile groups 246 security associations (SA) 311 security certificates 86 security policies about 187 defining 187 security profile groups, defining 246 security profiles about 211 actions 211 defining 246 security zones about 151 defining 151 in NAT policies 194 in security policies 188 interfaces 151 sensitive information, protecting 45 servers defining Kerberos 67 defining LDAP 66 defining RADIUS 65 defining syslog 76 service groups defining 240 service groups, defining 240 services, defining 239 session browser 267 shared gateways about 112 common interface 112 configuring 115 L3 interfaces 115 shared policy defining 382, 393 master device 387 Shortest Path Tree (SPT) 170 signatures custom 244 spyware 244 vulnerability 244 **SNMP** community string 47 defining trap destinations 75 MIB setup 47 MIBs 75 SNMP trap destinations defining 75 in logging profiles 247 software downgrading 53 upgrading 50, 61, 392, 417, 419 upgrading Panorama 419 version 252 source-specific multicast (SSM) 171 speed, link 128, 139, 149, 150 SSL decryption policies 246

decryption rule settings 200, 203 defining decryption policies 202 tech notes reference 202 SSL VPNs about 355 authentication profiles 62 comfort page 116 local user database 64 split tunnels 349 storage partitions 373 Panorama 373 sub category application 233 filtering 231 support information 117 support information, viewing 117 supported browsers 27 SYN flood 178 syslog servers custom syslog fields 77 defining 76 in logging profiles 247 system log about 69 viewing 267 system settings 115

Т

tables, using in web interface 26 tags on L2 subinterfaces 129, 140 on L3 subinterfaces 135 on virtual wires 126 tap interfaces, defining 149 tap mode deployment option 125 description 125 terminal services agent (TS Agent) 301 threat list 255 threat log 248 about 69 defining remote logging 246 viewing 266 threats ACC list 255 updating definitions 55 thresholds, alarm 178 time setting 30, 49 zone 31 traffic log 247 defining remote logging 246 viewing 266 Transport Layer Security (TLS) 66 TS Agent configuring firewall to support 301 configuring on the terminal server 303 installing 302

uninstalling 307 upgrading 302 tunnel interfaces 315 tunnel monitor fail over 177 profiles 177 wait-recover 177 tunneled traffic, and QoS 356 tunnels about VPN 311 number of IPSec 311 setting up 315 split for SSL VPNs 349 types of deployments 120 typographical conventions 15

U

UDP flood 179 unknown applications identifying 276 requesting App-ID 277 taking action 277 unnumbered loopback interfaces 146 upgrading Panorama software 417, 419 PAN-OS software 50, 61, 392 schedules 56 threat and application definitions 55 with high availability 51 URL filtering ACC page 254 continue and override response page 116, 434 defining profiles 217 dynamic categorization 217 list 254 override settings 44 profile settings 217 response pages 116 viewing log 266 viewing logs 69 user account lockout 61, 391 user database, SSL VPN 64 user interface navigation 26 User-ID Agent captive portal configuration 290 configuring firewall 286 configuring for Active Directory 298 for Active Directory, about 296 installing for Active Directory 297 overview 281 uninstalling and upgrading for Active Directory 301 username and password requirements 57

V

version, software 252 viewing devices 387

logs 264 session browser 267 session information 267 virtual firewall 363 virtual routers configuring 155, 157, 169 defining 155 multicast settings 169 next hop 157 routing protocols 153 runtime statistics 171 virtual systems about 110 and security zones 111 communications among 111 defining 110, 113, 115 defining multiple 113 enabling 31 enabling multiple 31 internal traffic flow 111 multiple 111 policies 110 security zones 110 shared gateway common interface 112 shared gateways 112 virtual wire 120 defining 125 interfaces 141 interfaces, configuring 138 **VLANs** interfaces, defining 143 L2 interfaces 143 **VM-Series** firewall installation 363 installing 366 limitations 364 overview 363 requirements 364 troubleshooting 370 VMware ESX(i) 372 VPN about 310 IPSec and IKE crypto profiles 312 large scale VPN and Dynamic Routing Protocols 333 large scale VPN deployment 324 large scale VPN overview 323 sample configuration 320 setting up tunnels 311 SSL, about 355 VPN tunnels about 311 IKE 311 manual security keys 311 securing 311 setting up 313, 315 vSphere 372 vulnerability protection profiles 215, 217

W

web interface committing changes 25 navigation 26 required fields 26 supported browsers 27 using 23 using tables 26 wildcard custom URL categories 242 patterns for allow and block lists 218 WildFire configure forwarding 424 configuring firewall settings 423 configuring settings on the portal 426 dashboard 425 log descriptions 425 setup tasks 423 subscription 422 supported file types 422 using the portal 425 viewing reports 427 WINS servers 172

Х

XML API 19

Ζ

zones defining 151 in NAT policies 194 in security policies 188 protection profiles 178, 358