



# **Palo Alto Networks Certified Network Security Engineer (PCNSE) Study Guide**

Jan 2023

---

# Table of Contents

<b>How to Use This Study Guide</b>	<b>2</b>
<b>About the PCNSE Exam</b>	<b>2</b>
Exam Format	2
How to Take This Exam	3
Disclaimer	3
<b>Audience and Qualifications</b>	<b>3</b>
Skills Required	3
<b>Recommended Training</b>	<b>3</b>
<b>Domain 1: Core Concepts</b>	<b>4</b>
<b>1.1 Identify how Palo Alto Networks products work together to improve PAN-OS services</b>	<b>4</b>
1.1.1 Security components	4
1.1.2 Firewall components	11
1.1.3 Panorama components	13
1.1.4 PAN-OS subscriptions and the features they enable	14
1.1.5 Plugin components	14
1.1.6 Heatmap and BPA reports	14
1.1.7 Artificial intelligence operations (AIOps)/Telemetry	16
1.1.8 IPv6	18
1.1.9 Internet of things (IoT)	18
1.1.10 References	18
<b>1.2 Determine and assess appropriate interfaces or zone types for various environments</b>	<b>20</b>
1.2.1 Layer 2 interfaces	20
1.2.2 Layer 3 interfaces	20
1.2.3 Virtual wire (vwire) interfaces	21
1.2.4 Tap interfaces	23
1.2.5 Subinterfaces	23
1.2.6 Tunnel interfaces	25
1.2.7 Aggregate interfaces	26
1.2.8 Loopback interfaces	26
1.2.9 Decrypt mirror interfaces	26
1.2.10 VLAN interfaces	26
1.2.11 References	27
<b>1.3 Identify decryption deployment strategies</b>	<b>27</b>
1.3.1 Risks and implications of enabling decryption	27
1.3.2 Use cases	29
1.3.3 Decryption types	29
1.3.4 Decryption profiles and certificates	30
1.3.5 Create a decryption policy in the firewall	31
1.3.6 Configure SSH Proxy	32
1.3.7 References	32

<b>1.4 Enforce User-ID</b>	<b>32</b>
1.4.1 Methods of building user-to-IP mappings	32
1.4.2 Determine if User-ID agent or agentless should be used	35
1.4.3 Compare and contrast User-ID agents	35
1.4.4 Methods of User-ID redistribution	35
1.4.5 Methods of group mapping	36
1.4.6 Server profile and authentication profile	37
1.4.7 References	37
<b>1.5 Determine how and when to use the Authentication policy</b>	<b>38</b>
1.5.1 Purpose of, and use case for, the Authentication policy	38
1.5.2 Dependencies	38
1.5.3 Captive portal versus GlobalProtect (GP) client	39
1.5.4 References	40
<b>1.6 Differentiate between the fundamental functions that reside on the management plane and data plane</b>	<b>40</b>
1.6.1 References	41
<b>1.7 Define multiple virtual systems (multi-vsys) environment</b>	<b>42</b>
1.7.1 User-ID hub	42
1.7.2 Inter-vsys routing	44
1.7.3 Service routes	47
1.7.4 References	48
<b>Domain 2: Deploy and Configure Core Components</b>	<b>48</b>
<b>2.1 Configure Management Profiles</b>	<b>48</b>
2.1.1 Interface Management Profile	49
2.1.2 SSL/TLS profile	50
2.1.3 References	50
<b>2.2 Deploy and configure Security Profiles</b>	<b>50</b>
2.2.1 Custom configuration of different Security Profiles and Security Profile Groups	51
2.2.2 Relationship between URL filtering and credential theft prevention	62
2.2.3 Use of username and domain name in HTTP header insertion	63
2.2.4 DNS Security	63
2.2.5 How to tune or add exceptions to a Security Profile	63
2.2.6 Compare and contrast threat prevention and advanced threat prevention	64
2.2.7 Compare and contrast URL Filtering and Advanced URL Filtering	66
2.2.8 References	67
<b>2.3 Configure zone protections, packet buffer protection, and DoS protection</b>	<b>68</b>
2.3.1 Customized values versus default settings	68
2.3.2 Classified versus aggregate profile values	72
2.3.3 Layer 3 and Layer 4 header inspection	73
2.3.4 References	74
<b>2.4 Design the deployment configuration of a Palo Alto Networks firewall</b>	<b>74</b>
2.4.1 Advanced high availability (HA) deployments	74

2.4.2 HA Pair	75
2.4.3 Zero-Touch Provisioning	75
2.4.4 Bootstrapping	75
2.4.5 References	77
<b>2.5 Configure authorization, authentication, and device access</b>	<b>77</b>
2.5.1 Role-based access control for authorization	77
2.5.2 Different methods used to authenticate	78
2.5.3 The Authentication Sequence	83
2.5.4 The device access method	83
2.5.5 References	83
<b>2.6 Configure and manage certificates</b>	<b>83</b>
2.6.1 Usage	83
2.6.2 Profiles	85
2.6.3 Chains	85
2.6.4 References	85
<b>2.7 Configure routing</b>	<b>86</b>
2.7.1 Dynamic routing	86
2.7.2 Redistribution Profiles	88
2.7.3 Static routes	89
2.7.4 Route monitoring	89
2.7.5 Policy-based forwarding	89
2.7.6 Virtual routers versus logical routers	89
2.7.7 References	92
<b>2.8 Configure NAT</b>	<b>93</b>
2.8.1 NAT policy rules	93
2.8.2 Security rules	94
2.8.3 Source NAT	94
2.8.4 No-NAT Policies	95
2.8.5 Use session browser to find NAT rule name	95
2.8.6 U-Turn NAT	96
2.8.7 Check HIT counts	97
2.8.7 References	99
<b>2.9 Configure site-to-site tunnels</b>	<b>99</b>
2.9.1 IPsec components	100
2.9.2 Static peers and dynamic peers for IPsec	100
2.9.3 IPsec tunnel Monitor Profiles	101
2.9.4 IPsec tunnel testing	101
2.9.5 Generic Routing Encapsulation	102
2.9.6 One-to-one and one-to-many tunnels	105
2.9.7 Determine when to use proxy IDs	106
2.9.8 References	107
<b>2.10 Configure service routes</b>	<b>107</b>



2.10.1 Default	108
2.10.2 Custom	110
2.10.3 Destination	110
2.10.4 Custom routes for different virtual systems versus destination routes	111
2.10.5 How to verify service routes	112
2.10.6 References	113
<b>2.11 Configure application-based QoS</b>	<b>113</b>
2.11.1 Enablement requirements	113
2.11.2 QoS policy rule	113
2.11.3 Add a Differentiated Services Code Point/ToS component	114
2.11.4 QoS Profile	115
2.11.5 Determine how to control bandwidth use on a per-application basis	115
2.11.6 Use QoS to monitor bandwidth utilization	115
2.11.6 References	116
<b>Domain 3: Deploy and Configure Features and Subscriptions</b>	<b>116</b>
<b>3.1 Configure App-ID</b>	<b>116</b>
3.1.1 Create security rules with App-ID	116
3.1.2 Convert port and protocol rules to App-ID rules	120
3.1.3 Identify the impact of application override to overall firewall functionality	122
3.1.4 Create custom apps and threats	124
3.1.5 Review App-ID dependencies	124
3.1.6 References	125
<b>3.2 Configure GlobalProtect</b>	<b>126</b>
3.2.1 GlobalProtect licensing	126
3.2.2 Configure the gateway and the portal	126
3.2.3 GlobalProtect agent	126
3.2.4 Differentiate between logon methods	127
3.2.5 Configure clientless VPN	127
3.2.6 HIP	127
3.2.7 Configure multiple gateway agent profiles	127
3.2.8 Split tunneling	128
3.2.9 References	128
<b>3.3 Configure decryption</b>	<b>129</b>
3.3.1 Inbound decryption	129
3.3.2 SSL forward proxy	129
3.3.3 SSL decryption exclusions	129
3.3.4 SSH proxy	130
3.3.5 References	130
<b>3.4 Configure User-ID</b>	<b>131</b>
3.4.1 User-ID agent and agentless	131
3.4.2 User-ID group mapping	131
3.4.3 Shared User-ID mapping across virtual systems	132

3.4.4 Data redistribution	132
3.4.5 User-ID methods	132
3.4.6 Benefits of using dynamic user groups (DUGs) in policy rules	133
3.4.7 Requirements to support dynamic user groups	136
3.4.8 How GlobalProtect internal and external gateways can be used	137
3.4.9 References	137
<b>3.5 Configure WildFire</b>	<b>137</b>
3.5.1 Submission profile	137
3.5.2 Action profile	138
3.5.3 Submissions and verdicts	138
3.5.4 Signature actions	139
3.5.5 File types and file sizes	141
3.5.6 Update schedule	141
3.5.7 Forwarding of decrypted traffic	142
3.5.8 References	142
<b>3.6 Configure Web Proxy</b>	<b>143</b>
3.6.1 Transparent proxy	143
3.6.2 Explicit proxy	144
3.6.3 References	144
<b>Domain 4: Deploy and Configure Firewalls Using Panorama</b>	<b>144</b>
<b>4.1 Configure templates and template stacks</b>	<b>144</b>
4.1.1 Components configured in a template	144
4.1.2 How the order of templates in a stack affects the configuration push to a firewall	144
4.1.3 Overriding a template value in a stack	144
4.1.4 Configure variables in templates	144
4.1.5 Relationship between Panorama and devices for dynamic update versions, policy implementation, and HA peers	144
4.1.6 References	145
<b>4.2 Configure device groups</b>	<b>145</b>
4.2.1 Device group hierarchies	145
4.2.2 Identify what device groups contain	146
4.2.3 Differentiate between different use cases for pre-rules, local rules, default rules, and post-rules	146
4.2.4 Identify the impact of configuring a primary device	147
4.2.5 Assign firewalls to device groups	147
4.2.6 References	149
<b>4.3 Manage firewall configurations within Panorama</b>	<b>149</b>
4.3.1 Licensing	149
4.3.2 Commit recovery feature	150
4.3.3 Automatic commit recovery	150
4.3.4 Commit types and schedules	151
4.3.5 Configuration backups	151

4.3.6 Commit type options	153
4.3.7 Manage dynamic updates for Panorama and Panorama-managed devices	154
4.3.8 Software and dynamic updates	155
4.3.9 Import firewall configurations into Panorama	155
4.3.10 Configure Log Collectors	155
4.3.11 Check firewall health and status from Panorama	156
4.3.12 Configure role-based access control on Panorama	156
4.3.13 References	157
<b>Domain 5: Manage and Operate</b>	<b>158</b>
<b>5.1 Manage and configure log forwarding</b>	<b>158</b>
5.1.1 Identify log types and criticalities	158
5.1.2 Manage external services	160
5.1.3 Create and manage tags	160
5.1.6 Log monitoring	160
5.1.4 Customize logging and reporting settings	160
5.1.5 References	161
<b>5.2 Plan and execute the process to upgrade a Palo Alto Networks system</b>	<b>161</b>
5.2.1 Single firewall	161
5.2.2 High availability pairs	162
5.2.3 Panorama push	162
5.2.4 Dynamic updates	163
5.2.5 References	164
<b>5.3 Manage HA functions</b>	<b>164</b>
5.3.1 Link monitoring	164
5.3.2 Path monitoring	166
5.3.3 HA links	168
5.3.4 Failover	173
5.3.5 Active/active and active/passive	173
5.3.6 HA interfaces	176
5.3.7 Clustering	177
5.3.8 Election setting	178
5.3.9 References	178
<b>Domain 6: Troubleshooting</b>	<b>179</b>
<b>6.1 Troubleshoot site-to-site tunnels</b>	<b>179</b>
6.1.1 IPSec	179
6.1.2 GRE	180
6.1.3 One-to-one and one-to-many tunnels	180
Phase 1 issues:	180
6.1.4 Route-based versus policy-based remote hosts	183
6.1.5 Tunnel monitoring	184
6.1.6 References	184
<b>6.2 Troubleshoot interfaces</b>	<b>185</b>

6.2.1 Transceivers	185
6.2.2 Settings	185
6.2.3 Aggregate interfaces, LACP	186
6.2.4 Counters	187
6.2.5 Tagging	187
6.2.6 References	188
<b>6.3 Troubleshoot Decryption</b>	<b>188</b>
6.3.1 Inbound decryption	188
6.3.2 SSL forward proxy	189
6.3.3 SSH proxy	190
6.3.4 Identify what cannot be decrypted and configure exclusions and bypasses	192
6.3.5 Certificates	192
6.3.6 References	193
<b>6.4 Troubleshoot routing</b>	<b>193</b>
6.4.1 Dynamic routing	193
6.4.2 Redistribution profiles	194
6.4.3 Static routes	194
6.4.4 Route monitoring	195
6.4.5 Policy-based forwarding	195
6.4.6 Multicast routing	195
6.4.7 Service routes	196
6.4.8 References	196
<b>6.5 General Troubleshooting</b>	<b>196</b>
6.5.1 Logs	196
6.5.2 Packet capture (pcap)	198
6.5.3 Reports	199
6.5.4 References	200
<b>6.6 Troubleshoot resource protections</b>	<b>200</b>
6.6.1 Zone Protection profiles	200
6.6.2 DoS protections	201
6.6.3 Packet buffer protections	201
6.6.4 References	203
<b>6.7 Troubleshoot GlobalProtect</b>	<b>203</b>
6.7.1 Portal and Gateway	203
6.7.2 Access to resources	205
6.7.3 GlobalProtect client	206
6.7.4 References:	206
<b>6.8 Troubleshoot policies</b>	<b>206</b>
6.8.1 NAT	206
6.8.2 Security	206
6.8.3 Decryption	207
6.8.4 Authentication	208

---

6.8.5 References	208
<b>6.9 Troubleshoot HA functions</b>	<b>209</b>
6.9.1 Monitor	209
6.9.2 Failover triggers	209
6.9.3 References	210
<b>Appendix D: Glossary</b>	<b>210</b>
<b>Continuing Your Learning Journey with Palo Alto Networks</b>	<b>220</b>

## How to Use This Study Guide

Welcome to the Palo Alto Networks PCNSE Study Guide. The purpose of this guide is to help you prepare for your PCNSE exam and achieve your PCNSE credential.

You can read through this study guide from start to finish, or you may jump straight to topics you would like to study. Hyperlinked cross-references will help you locate important definitions and background information from earlier sections.

## About the PCNSE Exam

The Palo Alto Networks Certified Network Security Engineer (PCNSE) is a formal, third-party proctored certification that indicates that those who have achieved it possess the in-depth knowledge to design, install, configure, maintain, and troubleshoot most implementations based on the Palo Alto Networks platform.

### Exam Format

**Certification name:** Palo Alto Networks Certified Network Security Engineer

**Delivered through Pearson VUE:** [www.pearsonvue.com/paloaltonetworks](http://www.pearsonvue.com/paloaltonetworks)

**Exam series:** PCNSE

**Seat time:** 80 minutes

**Number of items:** 70

**Format:** Multiple Choice, Scenarios with Graphics, and Matching

**Languages:** English and Japanese

Exam Domain	Weight (%)
Core Concepts	12%
Deploy and Configure Core Components	20%
Deploy and Configure Features and Subscriptions	17%
Deploy and Configure Firewalls Using Panorama	17%
Manage and Operate	16%
Troubleshooting	18%
<b>TOTAL</b>	<b>100%</b>

### How to Take This Exam

The exam is available through the third-party Pearson VUE testing platform. To register for the exam, visit: <https://home.pearsonvue.com/paloaltonetworks>

## Disclaimer

This study guide is intended to provide information about the objectives covered by this exam, related resources, and recommended courses. The material contained within this study guide is not intended to guarantee that a passing score will be achieved on the exam. Palo Alto Networks recommends that candidates thoroughly understand the objectives indicated in this guide and use the resources and courses recommended in this guide where needed to gain that understanding.

## Audience and Qualifications

The PCNSE exam should be taken by anyone who wants to demonstrate a deep understanding of Palo Alto Networks technologies, including customers who use Palo Alto Networks products, value-added resellers, pre-sales system engineers, system integrators, and support staff.

Candidates should have 3 to 5 years' experience working in the networking or security industries and the equivalent of 6 to 12 months' experience deploying and configuring Palo Alto Networks Next-Generation Firewalls within the Palo Alto Networks product portfolio.

## Skills Required

- You can plan, deploy, configure, operate, and troubleshoot Palo Alto Networks product portfolio components.
- You have product expertise and understand the unique aspects of the Palo Alto Networks product portfolio and how to deploy appropriately.
- You understand networking and security policies used by PAN-OS software.

## Recommended Training

Palo Alto Networks strongly recommends that you attend the following instructor-led training courses or equivalent virtual digital learning courses:

- Firewall Essentials: Configuration and Management (EDU-210) or digital learning
- Panorama: Managing Firewalls at Scale (EDU-220) or digital learning
- Firewall: Troubleshooting (EDU-330)

After you have completed the courses, practice on the platform to master the basics. Use the following resources to prepare for the exam. All resources can be found here:

<https://www.paloaltonetworks.com/services/education/certification.html#pcnse>

- Cybersecurity Skills Practice Lab
- PCNSE Study Guide and Practice Exam
- Administrator's Guide: Specific configuration information and "best practices"
- Preparation videos and tutorials

## Domain 1: Core Concepts

### 1.1 Identify how Palo Alto Networks products work together to improve PAN-OS services

#### 1.1.1 Security components

##### The Palo Alto Networks Cybersecurity Portfolio

The Palo Alto Networks cybersecurity portfolio is organized into three offerings: Strata for enterprise security, Prisma for cloud security, and Cortex for security operations. The following sections describe how they work together to address some of the world's greatest security challenges.

##### Strata: Enterprise Security

Strata prevents attacks with the industry-leading network security suite that enables organizations to embrace network transformation while consistently securing users, applications, and data—no matter where they reside. The suite brings together the following:

##### Machine Learning-Powered Next-Generation Firewalls

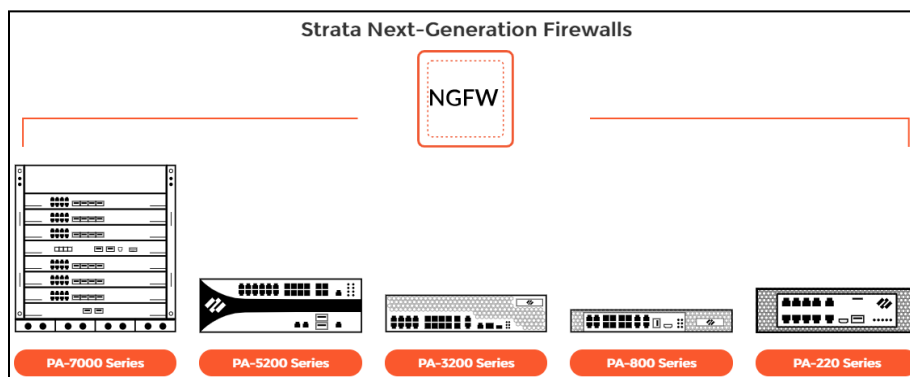
Palo Alto Networks Machine Learning (ML)-powered Next-Generation Firewalls (NGFWs) enable you to adopt best practices by using application, user, and content-based policies to minimize opportunities for attack. They are available as physical appliances, virtualized appliances, and cloud-delivered services—all of which can be managed with Panorama to ensure a consistent security stance.

The firewalls secure businesses with a prevention-focused architecture and with integrated innovations that are easy to deploy and use. Palo Alto Networks ML-powered NGFWs detect known and unknown threats, including those within encrypted traffic, using intelligence generated across thousands of customer deployments. The firewalls reduce risks and prevent a broad range of attacks. For example, they enable users to access data and applications based on business requirements and stop credential theft and an attacker's ability to use stolen credentials.

With these ML-powered NGFWs, it is quick and easy to create and maintain Security rules that mirror business policies and adapt to a dynamic environment. Automation reduces manual effort and accelerates response times with automated, policy-based actions and workflows that can be integrated with administrative tools, such as ticketing services via a RESTful API.

The family of NGFWs includes the following:

**PA-Series:** Various form factors of a PA-Series physical firewall can provide consistent protection for the entire network perimeter, from the headquarters, data center, and office campus to the branch offices and mobile and remote workforce. The models available include the PA-220, PA-800, PA-3200, PA-5200, and PA-7000 Series.





**VM-Series:** The virtualized version of the ML-powered NGFW offers the same level of protection as the PA-Series offerings. Further, it makes it easy to protect both private and public cloud deployments with segmentation and proactive threat prevention.



The VM-Series firewalls support the following virtualization environments:

- Amazon Web Services
- Cisco ACI
- Citrix NetScaler SDX
- Google CloudPlatform
- Kernel-based Virtual Machine (KVM)
- Microsoft Azure and Microsoft Hyper-V
- OpenStack
- VMware ESXi, VMware NSX, and VMware vCloud Air

### **Network Security Management: Panorama**

Panorama offers easy-to-implement and centralized management features to gain insights into network-wide traffic and threats and to administer NGFWs everywhere. Panorama is available in both appliance and virtual forms. Panorama provides the following:

- **Policy management**
- **Centralized visibility**
- **Network security insights**
- **Automated threat response**
- **Network security management**
- **Enterprise-level reporting and administration**

### **Prisma: Cloud Security**

Prisma Cloud delivers complete security across the development lifecycle on any cloud, enabling you to develop cloud-native applications with confidence. The Prisma suite includes Prisma Cloud, Prisma Access Secure Access Service Edge (SASE), Prisma SaaS, and the VM-Series ML-powered NGFWs.

### **Prisma Cloud**

Prisma Cloud is a Cloud Security Posture Management (CSPM) and cloud workload protection platform that provides comprehensive visibility and threat detection across an organization's hybrid, multi-cloud infrastructure.

Prisma Cloud taps into the cloud providers' APIs for read-only access to network traffic, user activity, and the configuration of systems and services. Prisma Cloud then correlates these disparate

datasets to help the cloud compliance and security analytics teams to prioritize risks and respond to issues quickly. Prisma Cloud also uses an agent-based approach to secure the host, container, and serverless computing environments against vulnerabilities, malware, and compliance violations.

The cloud-native security platform provides the following:

- **Comprehensive, cloud-native security**
- **Full lifecycle protection**
- **Protection across any cloud**

Prisma Cloud secures the following cloud-native infrastructures:

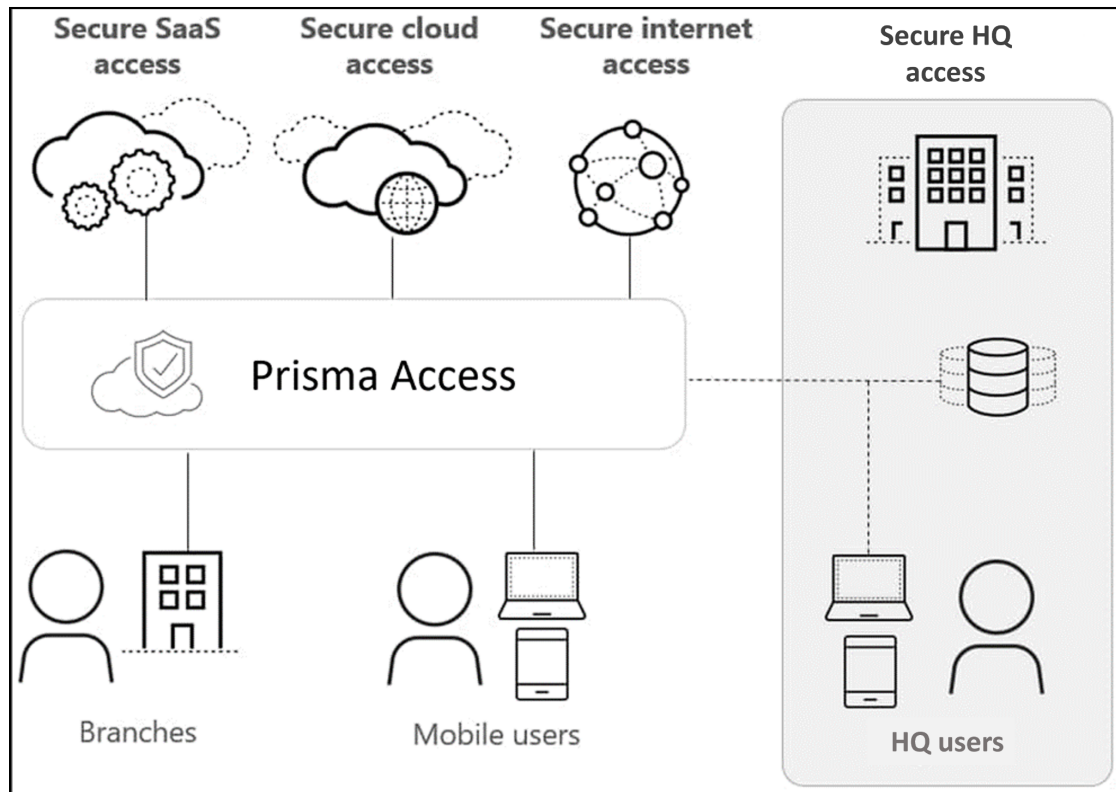
- Alibaba Cloud
- Amazon Web Services
- Docker EE
- Google CloudPlatform
- IBM Cloud
- Kubernetes
- Microsoft Azure
- Rancher
- Red Hat OpenShift
- VMware Tanzu

### **Prisma Access (SASE)**

Global expansion, mobile workforces, and cloud computing are changing the ways that organizations use to implement and deploy applications. Prisma Access provides the protection organizations need, where they need it. Prisma Access delivers a SASE (Secure Access Service Edge) that provides globally distributed networking and security to all users and applications in the organization.

SASE converges the capabilities of wide area networks (WANs) with network security to support the needs of the digital enterprise. These disparate networks and security services include software-defined wide area networks (SD-WANs), secure web gateways, cloud access security brokers (CASBs), software-defined perimeters, Domain Name System (DNS) protection, and firewall as a service.

Users connect to Prisma Access to access the internet and cloud and data center applications safely, regardless of their location.



### Prisma SaaS

Prisma SaaS (formerly known as Aperture) is a multimode CASB service that allows you to govern any sanctioned software as a service (SaaS) application use across all the users in your organization—and prevent risk from breaches and noncompliance. The service enables you to discover and classify data stored across supported SaaS applications, protect sensitive data from accidental exposure, identify and protect against known and unknown malware, and perform user activity monitoring to identify potential misuse or data exfiltration. Prisma SaaS delivers complete visibility and granular enforcement across all user, folder, and file activity within the sanctioned SaaS applications.

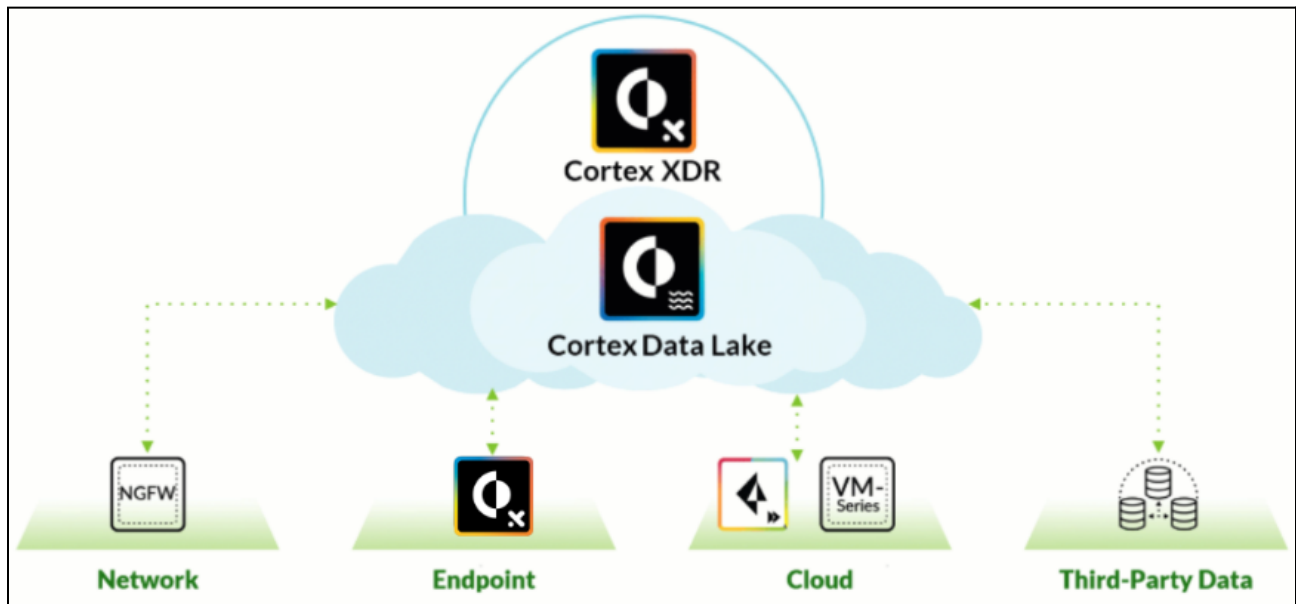
### Cortex: Security Operations

Cortex is the industry's most comprehensive product suite for security operations, empowering enterprises with best-in-class detection, investigation, automation, and response capabilities. The Cortex product suite includes Cortex XDR, Cortex XSOAR, Cortex Data Lake, and AutoFocus.

### Cortex XDR

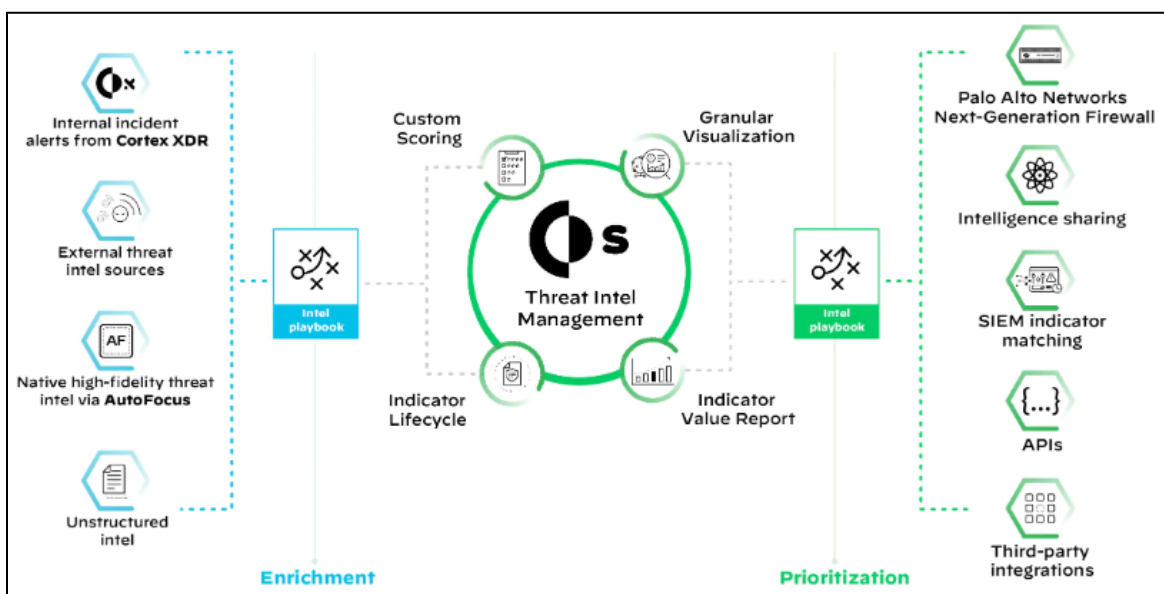
Cortex XDR is the industry's first extended detection and response platform that runs on integrated endpoint, network, and cloud data to reduce noise and focus on real threats. This platform provides complete visibility over network traffic, user behavior, and endpoint activity. It simplifies threat investigations by correlating logs from sensors to reveal threats and their timelines, which enables you to identify the root cause of every alert easily. Cortex XDR also allows you to perform immediate

response actions. Finally, to stop future attacks, you can proactively define indicators of compromise (IOCs) and behavioral indicators of compromise (BIOCs) to detect and respond to malicious activity. The following diagram depicts the Cortex XDR architecture.



### Cortex XSOAR

Cortex XSOAR is the industry's first extended Security Orchestration, Automation, and Response (SOAR) platform with native threat intelligence management. The SOAR technology can automate up to 95 percent of all of the response actions that require human review, thus allowing overloaded security teams to focus on more crucial tasks. Cortex XSOAR integrates with a wide variety of products to provide enhanced automation and response across processes. The following illustration depicts the Cortex XSOAR architecture, with its engine in the center, information sources on the left, and potential consumers on the right.

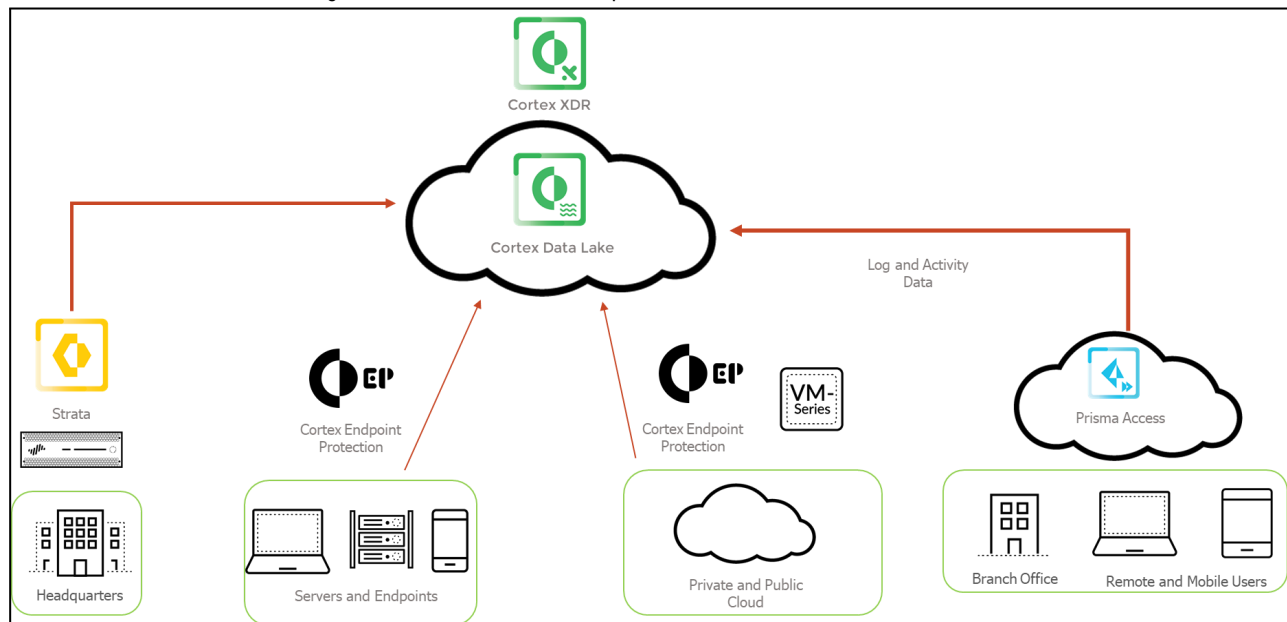


## Cortex Data Lake

Cortex Data Lake enables you to easily collect large volumes of log data so that innovative applications can gain insight from the organization's environment. You can simplify log infrastructure, automate log management, and use your data to prevent attacks more effectively. Cortex Data Lake can do the following:

- Radically simplify security operations by collecting, integrating, and normalizing the organization's security data
- Effortlessly run advanced artificial intelligence and ML with cloud-scale data
- Constantly learn from new data sources to evolve defenses

The following illustration depicts Cortex Data Lake as the central destination for information consolidation from many Palo Alto Networks products.



The following products can utilize Cortex Data Lake:

- Prisma Access
- Palo Alto Networks NGFWs and Panorama devices with the ability to connect to the cloud service
- Cortex XDR
- Previous versions of Palo Alto Networks Traps for endpoint protection and response (now Cortex XDR)
- Traps running version 5.0+ with the Traps management service

## 1.1.2 Firewall components

### Security Zones

Palo Alto Networks ML-powered NGFWs are zone-based. Zones designate a network segment in which all the nodes—users, data centers, demilitarized zone (DMZ) servers, and remote users—share similar network security requirements. The firewall security model is based on evaluating traffic as it passes from one zone to another. These zones act as a logical way to group physical and virtual interfaces. Zones are required to control and log the traffic that traverses the interfaces. All the defined interfaces should be assigned a zone that marks all of the traffic coming to or from the interface. Zones are defined for specific interface types—Tap, virtual wire, Layer 2, or Layer 3—and can be assigned to multiple interfaces of the same type only. An interface can be assigned only to a single zone, but a zone can contain multiple interfaces.

All sessions on the firewall are defined by source and destination zones. Rules can use these defined zones to allow or deny traffic, apply Quality of Service (QoS) policies, perform network address translation (NAT), and more. By default, all the traffic can flow freely within a zone and is referred to as intrazone traffic. Traffic between zones (called interzone traffic) is denied by default. Security policy rules are required to modify these default behaviors. Traffic can flow between zones only if a defined Security policy rule matches and allows the traffic. All the policies reference a source and destination zone, not interfaces, to match traffic. These policies are assessed in a top-down manner, meaning that the first rule with the appropriate match criteria will be matched.

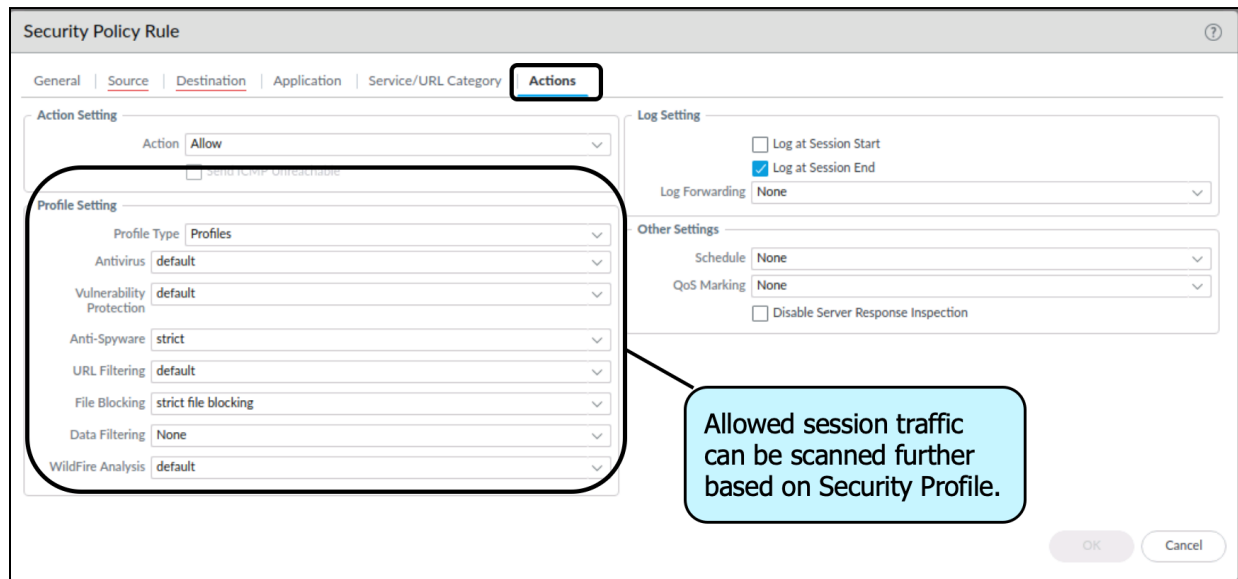
### Security Policy

Security policy rules are used to create a positive (allow list) and negative (block list) enforcement model for traffic flowing through the firewall. If logging is enabled on the matching policy, the action for that session is logged. These logs are extremely useful for adjusting the positive/negative enforcement model. The log information can be used to characterize traffic, thus providing specific use information and allowing precise policy creation and control. Log entries can be forwarded to external locations, including email and web servers, syslog servers, Panorama, and Cortex Data Lake.

Palo Alto Networks firewall logs, the ACC, App Scope, and other reporting tools all work to precisely describe traffic and use patterns.

You can use multiple match conditions to create these Security policy rules. Traffic-matching criteria can include security zones, source and destination IP addresses, and source and destination devices—as well as information about the application (App-ID), source user (User-ID), service (port), HIP match, and URL. The content of allowed sessions can be scanned based on Security Profiles (Content-ID) to identify unwanted traffic content. These profiles allow for the detection of both known and unknown threats through signatures and inline Machine-Learned models.

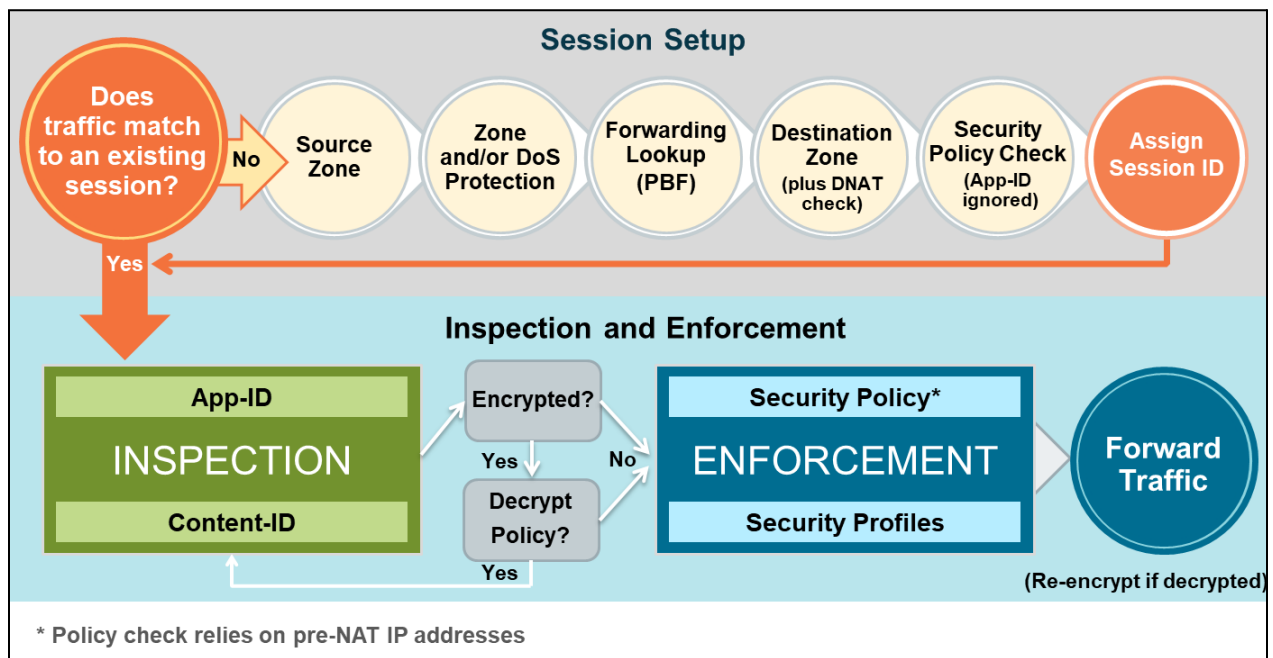
To utilize security profiles, simply add them to a Security policy rule as depicted below:



### Traffic Processing Sequence

The following image can help you visualize the Palo Alto Networks firewall processes. Understanding this traffic flow can help you better create an initial configuration, adjust the rules after installation, and troubleshoot existing rules.

Advanced analysis and discussion of the firewall flow logic is included in the Firewall: Troubleshooting (EDU-330) course.



The Palo Alto Networks NGFW was designed to use an efficient system known as next-generation processing. Next-generation processing enables packet evaluation, application identification, policy decisions, and content scanning in a single, efficient processing pass. This is known as Single Pass Parallel Processing architecture, or SP3.

Palo Alto Networks NGFWs contain the following next-generation features:

- **App-ID**
- **Content-ID**
- **User-ID**
- **Device-ID**

### 1.1.3 Panorama components

#### Panorama Overview


The PCNSE certification requires candidates to have knowledge of the Panorama firewall management functions. The following sections review these management concepts, but they do not cover all the Panorama features. Panorama offers several integration functions that provide enterprise management for multiple firewalls.

The Panorama management server provides centralized monitoring and management of multiple Palo Alto Networks NGFWs and Prisma Access deployments and of WildFire appliances and appliance clusters. The server provides a single location for overseeing all of the applications, users, and content traversing the network and, then, it uses this knowledge to create application enablement policies that protect and control the network. Panorama for centralized policy and firewall management increases operational efficiency in managing and maintaining a distributed network of firewalls.

Panorama uses device groups and templates to group firewalls into logical sets that require similar configuration. You can use the device groups and templates to manage all the configuration elements, policies, and objects on the managed firewalls centrally. Panorama also enables you to manage licenses, software (for example, PAN-OS software, SSL-VPN client software, GlobalProtect agent software), and content updates (for example, application and threat, WildFire, and antivirus updates) centrally.

Panorama's management web interface has the same look and feel as the firewall's management web interface.

Firewall menus available on the management web interface:



Panorama menus available on the management web interface:





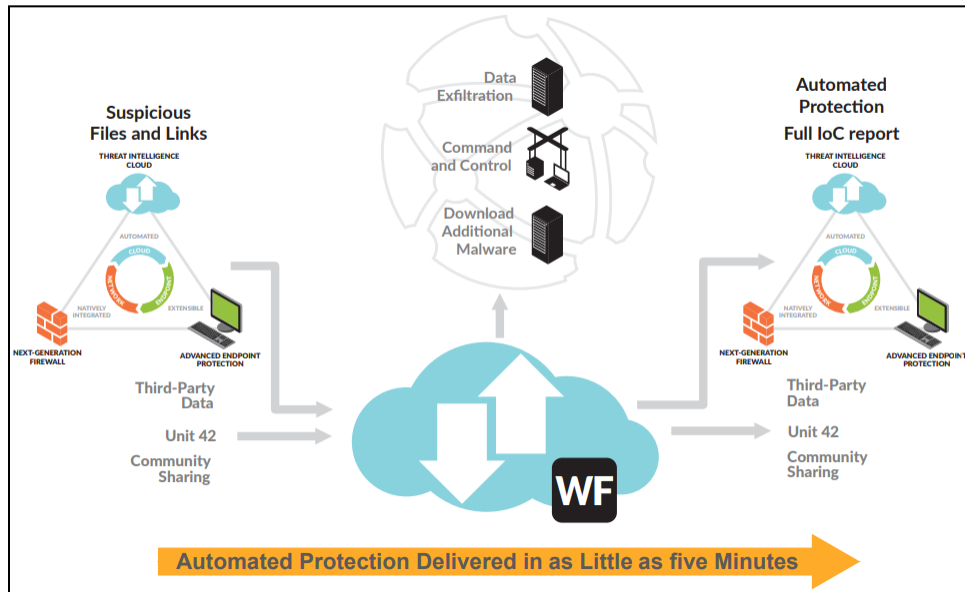
You can use the **Network** and **Device** tabs under “Templates” and the **Policies** and **Objects** tabs under “Device Groups” in Panorama to deploy a common base configuration for multiple firewalls with similar settings. To do this, you use a combination of a device group to manage shared policies and objects and a template stack (or multiple templates) to manage shared device and network settings.

### 1.1.4 PAN-OS subscriptions and the features they enable

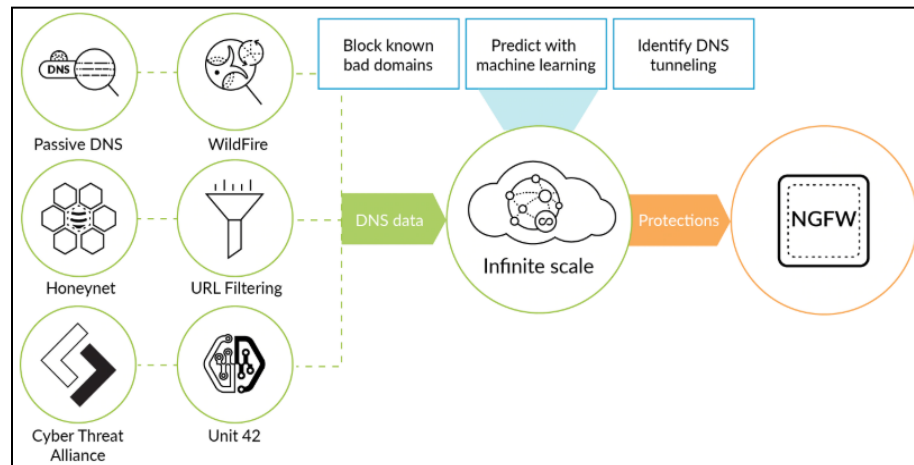
#### Security Subscriptions

Palo Alto Networks ML-powered NGFWs have a comprehensive range of security subscriptions natively integrated to provide comprehensive security that is automated and driven by ML. The subscriptions offered include the following:

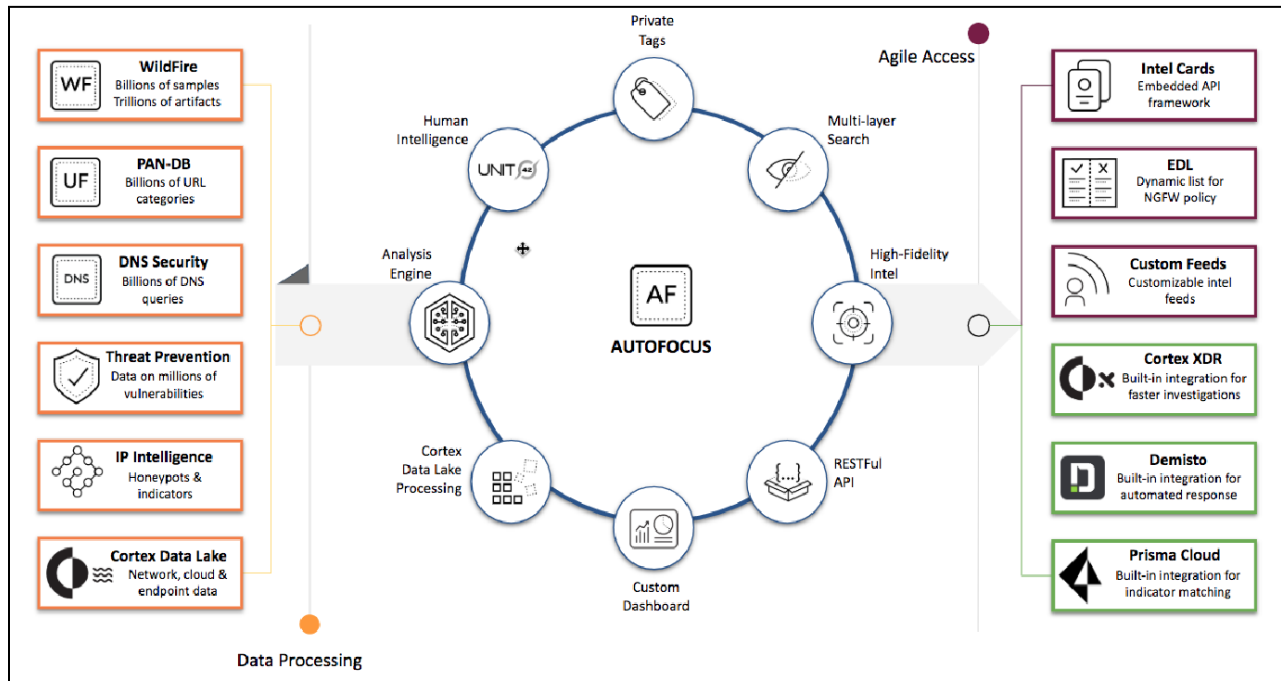
- **Threat Prevention**
- **Advanced URL Filtering**
- **WildFire**



- **DNS Security**



- **SD-WAN**
- **IoT Security**
- **AutoFocus**



### 1.1.5 Plugin components

Panorama plugins are available for both hardware and VM-Series panorama devices. Several plugins are available for Panorama, depending on the requirements. These plugins must be installed manually to extend Panorama’s native capabilities, such as managing the Prisma Access deployment.

The VM-Series plugin for Panorama is for managing VM-Series firewalls. It is preinstalled.

### 1.1.6 Heatmap and BPA reports

#### Comparing the Heatmap and BPA Reports

The free Best Practice Assessment (BPA) tool for Palo Alto Networks firewalls and Panorama evaluates a device’s configuration by measuring the adoption rate of a firewall’s capabilities and validating whether or not the policies adhere to best practices. The BPA tool provides recommendations and instructions about how to remediate the failed best practice checks. The goal for running the BPA tool is to reduce your attack surface. The BPA tool should be run on a scheduled basis (for example, quarterly) to ensure continuous improvement.

The two components of the BPA tool are the Security Policy Adoption Heatmap and the BPA assessment. The Heatmap analyzes a Palo Alto Networks deployment, measuring the adoption rate of features and capabilities across a targeted network infrastructure. The Heatmap can filter

information by device groups, serial numbers, zones, areas of architecture, and other categories. The results chart the progress of security improvement toward a Zero Trust network.

The BPA assessment compares a firewall or Panorama configuration against best practices and provides recommendations to strengthen the organization's security posture by fully adopting the Palo Alto Networks prevention capabilities. More than 200 security checks are performed on the firewall or Panorama configuration. A pass/fail score is provided for each check. If a check returns a failing score, the tool provides a justification for the failing score along with recommendations about how to resolve the issue.

Both components require the tech support file from either Panorama or a firewall to be uploaded to the Palo Alto Networks Customer Support Portal. After importing the tech support file, you should complete the architecture mapping, which maps existing zone names to predefined architecture classifications. Examples of architecture classifications are Enterprise – Perimeter – Internet, Internal – Core – Users, and Mobility – Remote Users/VPN.

### Heatmap Measurements

The Heatmap measures the adoption rate of Palo Alto Networks features. The results display the adoption rate based on source zone to destination zone. Column filters are available to allow you to examine specific device groups, source zones, and destination zones.

The Heatmap measures the adoption rate of the following Palo Alto Networks firewall features:

- WildFire
- Threat Prevention
- Anti-Spyware
- DNS Sinkhole
- Antivirus
- Vulnerability Protection
- URL Filtering
- File Blocking
- Data Filtering
- User-ID
- App-ID
- Service/Port
- Logging

### 1.1.7 Artificial intelligence operations (AIOps)/Telemetry

AIOps stands for 'artificial intelligence for IT operations'. It refers to the platforms that leverage ML and analytics to automate IT operations. AIOps harnesses big data from operational appliances and has the unique ability to detect and respond to issues instantaneously. Using the power of ML, AIOps strategizes using the various forms of data it compiles to yield automated insights that work to refine and iterate continually.

The first step in the process involves data extraction. Tools must collect the data coming from various systems and then cluster it in an appropriate manner, which makes the next step in the process most efficient. Then, a thorough analysis of the aggregated data is conducted. Using ML

algorithms, these tools detect patterns and relationships between pieces of data while identifying root problems and focal points within a system. In the next stage, AIOps looks to apply its “critical thinking skills” to react to the findings of the previous analysis. This entails deploying an automated optimization of IT operations while also using the detected patterns to learn and funnel closer to the potential pain points. This technology is generally paired with the ability to provide comprehensive analytical reports that help people make more intelligent, data-driven decisions.

When enabled, Telemetry allows the firewall to collect and forward traffic information to Palo Alto Networks. The data collected pertains to applications, threats, device health, and passive DNS information. All the Palo Alto Networks customers benefit from the data with improved accuracy and learning in threat findings and its community-driven approach in threat prevention. The data and source shared is maintained as anonymous and not shared with any external or third-party organizations.

### 1.1.8 IPv6

The firewall implementation of Neighbor Discovery (ND) is enhanced so that you can provision the IPv6 hosts with the Recursive DNS Server (RDNSS) option and DNS Search List (DNSSL) option, per RFC 6106, IPv6 Router Advertisement Options for DNS Configuration. When you configure Layer 3 interfaces, you configure these DNS options on the firewall so it can provision the IPv6 hosts; therefore, you don't need a separate DHCPv6 server to provision the hosts. The firewall sends IPv6 Router Advertisements (RAs) containing these options to the IPv6 hosts as part of their DNS configuration to fully provision them to reach internet services. Thus, the IPv6 hosts are configured with:

- The addresses of the RDNS servers that can resolve DNS queries
- A list of domain names (suffixes) that the DNS client appends (one at a time) to an unqualified domain name before entering the domain name into a DNS query

The IPv6 Router Advertisement for DNS configuration is supported for Ethernet interfaces, subinterfaces, Aggregated Ethernet interfaces, and Layer 3 VLAN interfaces on all of the PAN-OS platforms.

After you configure the firewall with the addresses of RDNS servers, the firewall provisions an IPv6 host (the DNS client) with those addresses. The IPv6 host uses one or more of those addresses to reach an RDNS server. Recursive DNS refers to a series of DNS requests by an RDNS server.

An IPv6 Router Advertisement can contain multiple DNS Recursive Server Address options, each with the same or different lifetimes. A single DNS Recursive DNS Server Address option can contain multiple Recursive DNS server addresses as long as the addresses have the same lifetime.

### 1.1.9 Internet of things (IoT)

The IoT Security solution works with next-generation firewalls to dynamically discover and maintain a real-time inventory of the IoT devices on your network. Through AI and ML algorithms, the IoT Security solution achieves a high level of accuracy, even classifying IoT device types encountered for the first time. IoT Security also provides the automatic generation of policy recommendations to control IoT device traffic, as well as the automatic creation of IoT device attributes for use in firewall policies.

## 1.1.10 References

- VM-Series Plugin and Panorama Plugins,  
<https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/panorama-plugins/plugins-types>
- What is AIOps,  
<https://www.paloaltonetworks.com/cyberpedia/what-is-aiops>
- Telemetry,  
<https://live.paloaltonetworks.com/t5/best-practice-assessment-device/telemetry-interpreting-bpa-checks/ta-p/336855#:~:text=Telemetry%20when%20enabled%20the%20firewall%20will%20collect%20and,applications%2C%20threats%2C%20device%20health%20and%20passive%20dns%20information.>
- IPv6 Router Advertisements for DNS Configuration,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/configure-interfaces/layer-3-interfaces/manage-ipv6-hosts-using-ndp/ipv6-router-advertisements-for-dns-configuration>
- IoT Security,  
<https://docs.paloaltonetworks.com/iot>

## 1.2 Determine and assess appropriate interfaces or zone types for various environments

### 1.2.1 Layer 2 interfaces

In a Layer 2 deployment, the firewall provides switching between two or more networks. Devices are connected to a Layer 2 segment; the firewall forwards the frames to the proper port, which is associated with the Media Access Control (MAC) address that is identified in the frame. A Layer 2 interface should be configured when switching is required. You must also configure the VLAN objects to associate with the Layer 2 interfaces. This is different from 802.1Q tagging. In addition to this, if routing for hosts on Layer 2 interfaces is required, you might need to create and associate a logical VLAN interface.

### 1.2.2 Layer 3 interfaces

In a Layer 3 deployment, the firewall routes traffic between multiple ports by using Transmission Control Protocol/Internet Protocol (TCP/IP) addressing. Before you can configure any Layer 3 interfaces, you must configure the virtual routers that you want the firewall to use to route the traffic for each Layer 3 interface.

Layer 3 deployments require more network planning and configuration preparation than most other firewall interfaces, but they remain the most widely used in firewall deployments. Palo Alto Networks supports both IPv4 and IPv6, simultaneously, through a dual-stack implementation when IPv6 is required.

Each Layer 3 interface must be configured with an IPv4 and/or an IPv6 address, zone assignment, and the attached virtual router that services the traffic on the interface. The options available to meet other connectivity requirements include the following:

- NetFlow integration
- Maximum segment size (MSS) adjustment
- Maximum transmission unit (MTU) adjustment
- Binding of firewall services (such as ping responses, web management interface availability)
- Neighbor discovery for IPv6
- Manual MAC address assignment
- Link Layer Discovery Protocol (LLDP) enablement
- Dynamic DNS support
- Link negotiation settings

### 1.2.3 Virtual wire (vwire) interfaces

In a virtual wire deployment, you install a firewall transparently on a network segment by binding two firewall ports (interfaces) together. The virtual wire logically connects the two interfaces; therefore, the virtual wire is internal to the firewall.

Use a virtual wire deployment only when you want to seamlessly integrate a firewall into a topology and when the two connected interfaces on the firewall do not need to perform any switching or routing. For these two interfaces, the firewall is considered a transparent bump in the wire.

A virtual wire deployment simplifies firewall installation and configuration because you can insert the firewall into an existing topology without assigning MAC or IP addresses to the interfaces, redesigning the network, or reconfiguring the surrounding network devices. The virtual wire supports the blocking or allowing of traffic based on the virtual LAN (VLAN) tags. The virtual wire also supports Security policy rules, App-ID, Content-ID, User-ID, decryption, LLDP, active/passive and active/active high availability (HA), QoS, zone protection (with some exceptions), non-IP protocol protection, denial of service (DoS) protection, packet buffer protection, tunnel content inspection, and NAT.

Each virtual wire interface is directly connected to a Layer 2 or Layer 3 networking device or host. The virtual wire interfaces have no Layer 2 or Layer 3 addresses. When a virtual wire interface receives a frame or packet, it ignores any Layer 2 or Layer 3 addresses for switching or routing purposes; however, it applies the organization's security or NAT policy rules before passing an allowed frame or packet over the virtual wire to the second interface and on to the network device connected to it.

You would not use a virtual wire deployment for the interfaces that need to support switching, VPN tunnels, or routing because they require a Layer 2 or Layer 3 address. A virtual wire interface does not use an Interface Management Profile.

All the firewalls that are shipped from the factory have two Ethernet ports (port 1 and port 2) preconfigured as virtual wire interfaces. These interfaces allow all of the untagged traffic.

## 1.2.4 Tap interfaces

A network tap is a device that provides a way to access the data that is flowing across a computer network. TAP mode deployments allow you to passively monitor traffic flows across a network by using a switch port analyzer (SPAN) or mirror port.

A switch's SPAN or mirror port permits the copying of traffic from ports on the switch to the tap interface of the firewall, providing a one-way flow of copied network traffic into the firewall. This configuration allows the firewall to detect traffic and threats but prevents any enforcement action because the traffic does not flow through the firewall back to the environment.

By utilizing TAP interfaces, you can get visibility into which applications are running on the network as well as identify threats without having to make any changes to your network design. Remember, however, that the traffic is not running through the firewall when it arrives via a TAP interface; so no action can be taken on the traffic, including blocking any traffic that includes threats or applying QoS traffic control.

## 1.2.5 Subinterfaces


Layer 3, Layer 2, and VWIRE interfaces can all utilize subinterfaces to separate traffic into different zones. Using the VLAN tags (802.1Q) to differentiate and classify traffic, we gain flexibility and can apply relevant policies according to environmental requirements. VWIRE subinterfaces also allow the use of IP classifiers to further this segregation.

**VLAN tags in conjunction with IP classifiers (address, range, or subnet):** The following example shows an 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 how IP classifiers are used to classify traffic into separate zones and apply relevant policy for customers from each network.

### VIRTUAL WIRE SUBINTERFACE WORKFLOW

- Configure two Ethernet interfaces as type virtual wire. Assign these interfaces to a virtual wire.
- Create subinterfaces on the parent virtual wire to separate Customer A and Customer B traffic. Make sure that the VLAN tags defined on each pair of subinterfaces configured as virtual wire(s) are identical. This is essential because a virtual wire does not switch VLAN tags.
- Create new subinterfaces and define IP classifiers. This task is optional and only required if you want 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.

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

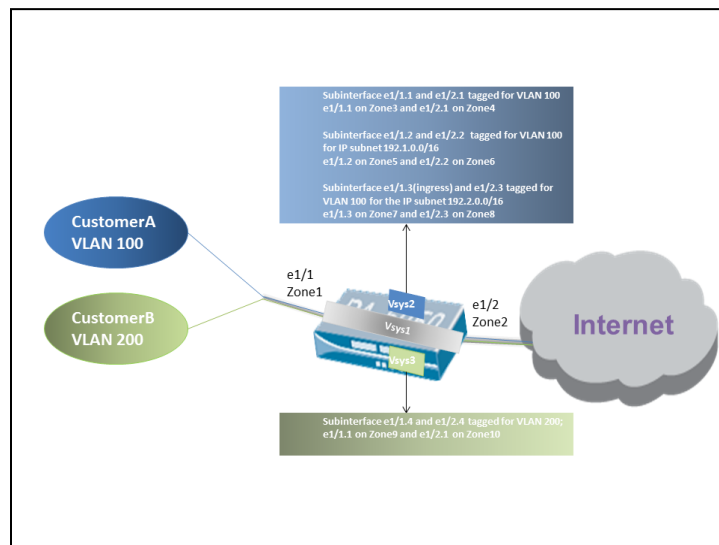


Virtual wire deployment with subinterfaces (VLAN tags only) depicts Customer A and Customer B connected to the firewall through one physical interface, ethernet1/1, configured as a virtual wire; it is the ingress interface. A second physical interface, ethernet1/2, is also part of the virtual wire; it is the egress interface that provides access to the internet.

For Customer A, you also have the subinterfaces ethernet1/1.1 (ingress) and ethernet1/2.1 (egress). For Customer B, you have the subinterfaces ethernet1/1.2 (ingress) and ethernet1/2.2 (egress). When configuring the subinterfaces, you must assign the appropriate VLAN tag and zone to apply policies for each customer. In this example, the policies for Customer A are created between Zone1 and Zone2 while the policies for Customer B are created between Zone3 and Zone4.

When traffic enters the firewall from Customer A or Customer B, 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; therefore, that subinterface is selected. The policies defined for the zone are evaluated and applied before the packet exits from the corresponding subinterface.

Virtual wire deployment with subinterfaces (VLAN Tags and IP Classifiers) depicts Customer A and Customer B 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.





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 of the tagged and untagged traffic except VLAN tags 100 and 200, which are assigned to the subinterfaces.

Customer A is managed on vsys2, and Customer B is managed on vsys3. On vsys2 and vsys3, the following virtual wire subinterfaces are created with the appropriate VLAN tags and zones to enforce policy measures.

CUSTOMER	VSYS	VIRTUAL WIRE SUBINTERFACES	ZONE	VLAN TAG	IP CLASSIFIER
A	2	e1/1.1 (ingress) e1/2.1 (egress)	Zone 3 Zone 4	100 100	None
	2	e1/1.2 (ingress) e1/2.2 (egress)	Zone 5 Zone 6	100 100	IP subnet 192.1.0.0/16
	2	e1/1.3 (ingress) e1/2.3 (egress)	Zone 7 Zone 8	100 100	IP subnet 192.2.0.0/16
B	3	e1/1.4 (ingress) e1/2.4 (egress)	Zone 9 Zone 10	200 200	None

When traffic enters the firewall from Customer A or Customer B, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this case, for Customer A, multiple subinterfaces use the same VLAN tag. Therefore, 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. It then selects the appropriate virtual wire to route traffic through the accurate subinterface.

### 1.2.6 Tunnel interfaces

In a VPN tunnel setup, the Layer 3 interface at each end must have a logical tunnel interface for the firewall to connect to and establish a VPN tunnel. A tunnel interface is a logical (virtual) interface that is used to deliver traffic between two endpoints. If you configure a proxy ID, the proxy ID is counted toward any Internet protocol security (IPsec) tunnel capacity.

The tunnel interface must belong to a security zone to apply policy, and it must be assigned to a virtual router to use the existing routing infrastructure. Ensure that the tunnel interface and the physical interface are assigned to the same virtual router so that the firewall can perform a route lookup and determine the appropriate tunnel to use.

The Layer 3 interface that the tunnel interface is attached to typically belongs to an external zone—for example, the untrust zone. Although the tunnel interface can be in the same security zone as the physical interface, you can create a separate zone for the tunnel interface for added security and better visibility. If you create a separate zone for the tunnel interface, such as a VPN

zone, you will need to create Security policies to allow traffic to flow between the VPN zone and the trust zone.

A tunnel interface does not require an IP address to route traffic between the sites. An IP address is required only if you want to enable tunnel monitoring or if you are using a dynamic routing protocol to route traffic across the tunnel. With dynamic routing, the tunnel IP address serves as the next hop IP address for routing traffic to the VPN tunnel.

If you are configuring the Palo Alto Networks firewall with a VPN peer that performs policy-based VPN, you must configure a local and remote proxy ID when setting up the IPsec tunnel. Each peer compares the proxy IDs that are configured on it with what is received in the packet to allow a successful IKE Phase 2 negotiation. If multiple tunnels are required, configure unique proxy IDs for each tunnel interface; a tunnel interface can have a maximum of 250 proxy IDs. Each proxy ID counts toward the IPsec VPN tunnel capacity of the firewall, and the tunnel capacity varies by the firewall model.

### 1.2.7 Aggregate interfaces

An Aggregate Ethernet (AE) interface group uses IEEE 802.1AX link aggregation to combine multiple Ethernet interfaces into a single virtual interface that connects the firewall to another network device or firewall. An AE interface group increases the bandwidth between peers by load balancing traffic across the combined interfaces. It also provides redundancy: When one interface fails, the remaining interfaces continue to support traffic.

Before you configure an AE interface group, you must configure its interfaces. Hardware media can differ among the interfaces assigned to an aggregate group. For example, you can mix fiber optic and copper but the bandwidth (1Gbps, 10Gbps, 40Gbps, or 100GBps) and interface type (HA3, virtual wire, Layer 2, or Layer 3) must be the same. You can add at least eight AE interface groups per firewall, although some firewall models support 16, and each group can have up to eight interfaces.

Aggregate interface creation begins with the definition of an Aggregate Interface group, after which individual interfaces are added to the group.

### 1.2.8 Loopback interfaces

Loopback interfaces are logical Layer 3 interfaces that exist only virtually and connect to the virtual routers in the firewall. Loopback interfaces are used for multiple network engineering and implementation purposes. They can be destination configurations for DNS sinkholes, GlobalProtect service interfaces (such as portals and gateways), routing identification, and more.

### 1.2.9 Decrypt mirror interfaces

Decryption mirroring is a special configuration that supports the routing of copied decrypted traffic through an external interface to another system, such as a data loss prevention (DLP) service. The Palo Alto Networks firewalls can automatically send a copy of decrypted traffic to a specified interface by using the Decryption Mirroring feature. This option can be licensed at no cost for midrange and high-end firewalls that can automatically forward copies of decrypted traffic to other devices.

## 1.2.10 VLAN interfaces

A VLAN interface can provide routing into a Layer 3 network (IPv4 and IPv6) from a Layer 2 deployment. You can add one or more Layer 2 Ethernet ports to a VLAN object and associate a VLAN object with a VLAN interface.

## 1.2.11 References

- Virtual Wire Interfaces, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/configure-interfaces/virtual-wire-interfaces>
- Configure Layer 3 Interfaces, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/configure-interfaces/layer-3-interfaces/configure-layer-3-interfaces>
- Tap Interfaces, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/configure-interfaces/tap-interfaces>
- Aggregate Ethernet (AE) Interface Group, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/network/network-interfaces/aggregate-ethernet-ae-interface-group>
- How to Configure a Decrypt Mirror Port, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10q000000CIGDCA0>
- Decryption Mirroring, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/decryption-concepts/decryption-mirroring>
- Network > Interfaces > VLAN, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/network/network-interfaces-vlan>

## 1.3 Identify decryption deployment strategies


### 1.3.1 Risks and implications of enabling decryption

#### Packet Visibility

The use of encryption in all of the network applications is growing rapidly. When traffic is encrypted, the Palo Alto Networks firewall loses visibility into the packet contents, thus making Content-ID impossible. Due to this lack of visibility, malware might be able to pass unchallenged to an endpoint, at which point it is decrypted and able to attack. Decryption policies maximize the firewall's visibility into packet content to allow for content inspection. In addition to this, decryption also allows the App-ID engine to identify applications at a more granular level.

#### Decryption

The Secure Sockets Layer (SSL) and Secure Shell (SSH) encryption protocols secure traffic between two entities, such as a web server and a client. SSL encapsulates traffic and encrypts data so that the data is meaningless to anyone other than the client and server with the correct certificates and keys to decode the data.



The Palo Alto Networks firewall decryption is policy-based and can be used to decrypt, inspect, and control both inbound and outbound SSL and SSH connections. Decryption policies enable you to specify traffic for decryption according to destination, source, user/user group, or URL category and to block or restrict the specified traffic according to your security requirements. The firewall uses certificates and keys to decrypt the traffic specified by the policy into plaintext. Once this is done, App-ID and Content-ID security analysis is performed on the plaintext traffic, including Antivirus, Vulnerability-Protection, Anti-Spyware, URL Filtering, Data Filtering, and File Blocking.

After traffic is decrypted and inspected on the firewall, the plaintext traffic is re-encrypted as it exits the firewall to ensure privacy and security.

### **Keys and Certificates**

Central to this discussion is the role of digital certificates to secure the SSL-encrypted data. Your understanding of this role and planning for proper certificate needs and deployment are important considerations in decryption use.

Encryption technology uses keys to transform cleartext strings into ciphertext. These keys are generated using passwords and other shared secrets. The Palo Alto Networks firewalls decrypt encrypted traffic by using keys to transform encrypted strings into cleartext. Certificates are used to establish the firewall as a trusted third party and to create a secure connection. SSL decryption (both forward proxy and inbound inspection) requires certificates to establish trust between two entities to secure an SSL/TLS connection. Certificates can also be verified when traffic is excluded from SSL decryption. You can integrate a hardware security module (HSM) with a firewall to enable enhanced security for the private keys used in the decryption of SSL Forward Proxy and SSL Inbound Inspection.

The use of certificates is central to other firewall functions in addition to decryption. This led to the implementation of extensive certificate management capabilities on the firewall.

### **App-ID and Encryption**

The App-ID scanning engine's effectiveness is often compromised by the encrypted traffic that prevents the scanning of packet contents for identifying elements. This traffic is typically given the App-ID of "SSL". In some cases, the App-ID engine can evaluate elements of the traffic, allowing the App-ID engine to assign App-IDs properly without scanning contents; however, this is not the case in a large number of situations. Decrypting the SSL traffic gives deeper visibility and control, exposing more granular applications within a given flow of traffic.

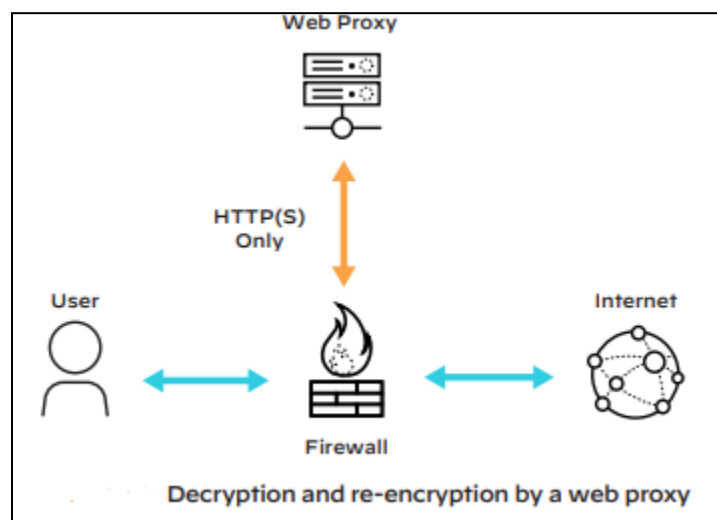
### 1.3.2 Use cases

Many technical options are available to decrypt traffic on the network, including web proxies, application delivery controllers, SSL visibility and decryption appliances, and NGFWs. Where it's best to decrypt TLS/SSL traffic depends on which option provides the greatest protection with the least management overhead.



### Web Proxies

A web proxy acts as a “middleman” by decrypting and inspecting outbound traffic before re-encrypting it and sending it to its destination (see the following figure). However, web proxies are limited to inspecting and securing web traffic, which includes HTTP and HTTPS. They are typically deployed on well-known web ports, such as 80 and 443. If an application uses non-web ports or protocols, web proxies can't see the traffic. For example, the Office/Microsoft 365 applications work across multiple ports besides 80/443. Regular proxies would miss traffic on these other ports. Moreover, web proxies cannot access non-web traffic, defeating the purpose of gaining complete visibility and control over encrypted traffic on a network. It would be like deploying airport security in only one major terminal and leaving the other terminals exposed. Proxies also require you to modify your browser's proxy settings or using a proxy autoconfig file, which adds more management overhead and an additional area to diagnose if users can't access the internet.

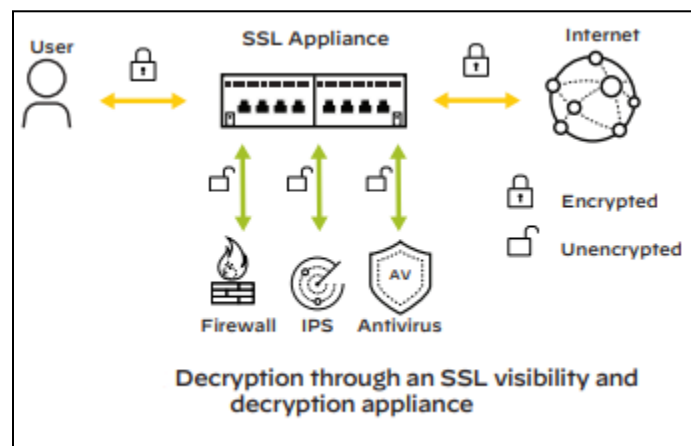


## Application Delivery Controllers

SSL offload is one of the functions performed by the Application Delivery Controllers (ADCs). An ADC deployment usually requires two separate appliances, one to decrypt traffic and the other to re-encrypt traffic. Given the two-stage operation, the problem with ADC deployments is that, once decrypted, the traffic travels unencrypted between the ADC devices until it hits the encryption device. An adversary can simply sniff the traffic and retrieve sensitive data in cleartext or manipulate the traffic. This undermines one of the fundamental purposes of encryption—the promise of complete confidentiality—and might violate compliance laws in some industries and geographies.

## SSL Visibility and Decryption Appliances

SSL visibility appliances decrypt traffic and make it available to all the other network security functions that need to inspect it, including web proxies, DLP systems, and antivirus programs. Managing and maintaining the links between these third-party devices and a decryption appliance also adds to the complexity and overhead.



### 1.3.3 Decryption types

The firewall provides three types of decryption policy rules: SSL Forward Proxy to control outbound SSL traffic, SSL Inbound Inspection to control inbound SSL traffic, and SSH Proxy to control tunneled SSH traffic. You can attach a Decryption Profile to a policy rule to apply granular access settings to traffic, including checks for server certificates, unsupported modes, and failures.

SSL decryption (both forward proxy and inbound inspection) requires certificates to establish the firewall as a trusted third party and to establish trust between a client and a server to secure an SSL/TLS connection. You can also use certificates when excluding servers from SSL decryption for technical reasons (the site breaks decryption for reasons, such as certificate pinning, unsupported ciphers, or mutual authentication). SSH decryption does not require certificates.

Note that through granular decryption policies, you can also add exceptions for the websites you do not want to decrypt due to legal and ethical reasons; for example, websites categorized as Health and Medicine, Finance, and Shopping.

### 1.3.4 Decryption profiles and certificates

A Decryption Profile controls SSL protocols, certificate verification, and failure checks to prevent traffic that uses weak algorithms or unsupported modes from accessing the network. We can use decryption profiles to dictate not only when not to decrypt, but also when to deny traffic based on weak encryption parameters. For example, SSL 3.0 is no longer considered secure. We can use a decryption profile to restrict SSL traffic to TLS 1.1 and above, even if we choose not to decrypt it.

When you apply a decryption policy with a restrictive decryption profile to traffic, a session between the client and the server is only established if the firewall trusts the certificate authority (CA) that signed the server certificate. To establish trust, the firewall must have the server root CA certificate in its certificate trust list (CTL) and use the public key contained in that root CA certificate to verify the signature. The firewall then presents a copy of the server certificate signed by the forward trust certificate for the client to authenticate. You can also configure the firewall to use an enterprise CA as a forward trust certificate for SSL Forward Proxy. If the firewall does not have the server root CA certificate in its CTL and if the decryption profile permits, the firewall presents a copy of the server certificate signed by the forward untrust certificate to the client. The forward untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites hosted by a server with untrusted certificates.

The following table describes the certificates that the Palo Alto Networks firewalls use for decryption:

CERTIFICATES USED WITH DECRYPTION	DESCRIPTION
Forward Trust (used for SSL Forward Proxy decryption)	The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate signed by a CA that the firewall trusts. By default, the firewall determines the key size to use for the client certificate based on the key size of the destination server. However, you can configure the key size for SSL proxy server certificates. For added security, consider storing the private key associated with the forward trust certificate on a hardware security module.
Forward Untrust (used for SSL Forward Proxy decryption)	The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall does not trust.
SSL Inbound Inspection	The certificates of the servers on your network for which you want to perform SSL Inbound Inspection of traffic destined for those servers. Import the server certificates onto the firewall.

### 1.3.5 Create a decryption policy in the firewall

You can configure the firewall to decrypt traffic for visibility, control, and granular security. Decryption policies can apply to SSL, including SSL encapsulated protocols such as Internet Message Access Protocol, Post Office Protocol Version 3, Simple Mail Transfer Protocol, File Transfer Protocol Secure, and Secure Shell [IMAP(S), POP3(S), SMTP(S), and FTP(S) and SSH] traffic. SSH decryption can be used to decrypt outbound and inbound SSH traffic to ensure that secure protocols are not being used to tunnel disallowed applications and content. Add a decryption policy

rule to define any traffic that you want to decrypt (for example, you can decrypt traffic based on URL categorization). Decryption policy rules are compared against the traffic in sequence, so more specific rules must precede the more general ones. SSL forward proxy decryption requires the configuration of a trusted certificate that is presented to the user if the server to which the user is connecting possesses a certificate signed by a CA trusted by the firewall. Create a certificate on the Device Certificate Management Certificates page, click the name of the certificate, and select **Forward Trust Certificate**.

Create a decryption policy rule to define traffic for the firewall to decrypt and the type of decryption you want the firewall to perform: SSL Forward Proxy, SSL Inbound Inspection, or SSH Proxy decryption. You can also use a decryption policy rule to define Decryption Mirroring.

### 1.3.6 Configure SSH Proxy

Configuring SSH proxy does not require certificates. The key used to decrypt SSH sessions is generated automatically on the firewall during bootup. With SSH decryption enabled, the firewall decrypts SSH traffic and blocks or restricts it based on your decryption policy and Decryption Profile settings. The traffic is re-encrypted as it exits the firewall.

### 1.3.7 References

- Keys and Certificates for Decryption Policies, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/decryption-concepts/keys-and-certificates-for-decryption-policies>
- Keys and Certificates, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/certificate-management/keys-and-certificates>
- How Palo Alto Networks Identifies HTTPS Applications Without Decryption, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIVSCA0>
- Decryption Exclusions, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/decryption-exclusions>
- Decryption Overview, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/decryption-overview>
- Decryption: Why, Where, and How, [https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en\\_US/resources/whitepapers/decryption-why-where-and-how](https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en_US/resources/whitepapers/decryption-why-where-and-how)



## 1.4 Enforce User-ID

### 1.4.1 Methods of building user-to-IP mappings

#### User-ID and Mapping Users

The User-ID feature of the Palo Alto Networks NGFW enables you to create policy rules and perform reporting based on users and groups rather than on individual IP addresses.

User-ID seamlessly integrates the Palo Alto Networks firewalls with a range of enterprise directory and terminal services offerings, thus enabling you to associate application activity and policy rules with users and groups, not just IP addresses. Furthermore, with User-ID enabled, the ACC, App Scope, reports, and logs can all include usernames in addition to user IP addresses.

For user-based and group-based policies, the firewall requires a list of all the available users and their corresponding group mappings that you can select when defining policies. The firewall collects group mapping information by connecting directly to the LDAP directory server. No other types of directory services are supported for group mapping.

Before the firewall can enforce user-based and group-based policies, it must be able to map the IP addresses (based on the packets it receives) to usernames. User-ID provides many mechanisms to collect this user-based mapping information.

A User-ID agent process runs either on the firewall (agentless implementation) or is installed as a separate process on a Microsoft Windows Server-based host. This User-ID agent monitors various network technologies for authentication events and gathers the data, creating a master IP-address-to-user mapping table stored in the firewall. For example, the User-ID agent monitors server logs for login events, probes clients, and listens for syslog messages from authenticating services. To identify mappings for IP addresses that the agent did not map, you can configure the firewall to redirect HTTP requests to a Captive Portal login. You can customize multiple user mapping mechanisms to suit your environment and even use different mechanisms at different sites.

#### Mapping IP Addresses to Usernames

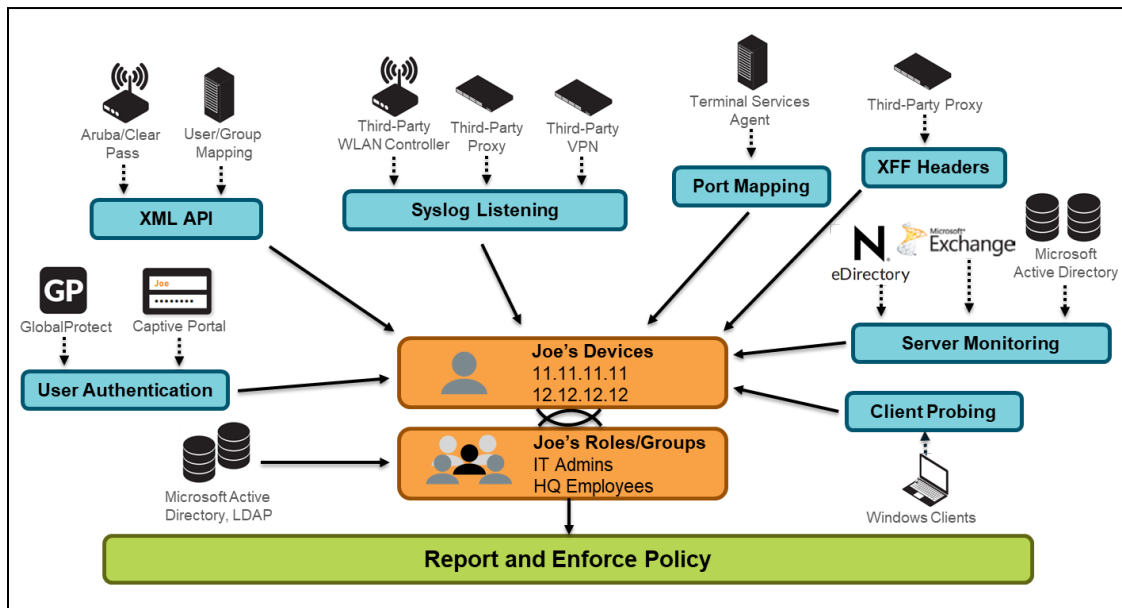
A user's IP address constantly changes because they use many devices and also because laptops provide unlimited mobility. Capturing that information regularly is difficult. The firewall needs to be able to simultaneously monitor multiple sources.

The firewall also has many ways to capture user information. The firewall can use server monitoring to monitor the Security logs on a Windows server for successful authentication events. Syslog monitoring of login events can be used with LDAP and Linux, among others.

The different methods of user mapping include:

- **Server monitoring:** A Windows-based User-ID agent, or the built-in PAN-OS integrated User-ID agent inside the PAN-OS firewall, monitors Security Event logs for successful login and logout events on Microsoft domain controllers, Exchange servers, or Novell eDirectory servers.
- **Port mapping:** For Microsoft Terminal Services or Citrix environments, users might share the same IP address. To overcome this issue, the Palo Alto Networks Terminal Services agent must be installed on the Windows or Citrix terminal server. The Terminal Services agent uses the source port of each client connection to map each user to a session. Linux terminal servers do not support the Terminal Services agent and must use the XML API to send user mapping information from login or logout events to User-ID.
- **Syslog:** The Windows-based User-ID agent and the PAN-OS integrated User-ID agent both use Syslog Parse Profiles to interpret login and logout event messages that are sent to syslog servers from the devices that authenticate users. Such devices include wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, and other network access control devices.
- **XFF headers:** If a proxy server exists between users and a firewall, the firewall might see the source IP address of the proxy server instead of the original source IP address of the host that originated the traffic. Most proxy servers have a feature that allows the forwarding of the original source IP address of the host to the firewall within an XFF header. Use of the original client source IP address enables the firewall to map the IP address to a username.
- **Authentication policy and Captive Portal:** The User-ID agent sometimes cannot map an IP address to a username by using any of the methods described. In these cases, you can use an Authentication policy and Captive Portal, whereby any web traffic (HTTP or HTTPS) that matches an Authentication policy rule forces the user to authenticate via one of the following three Captive Portal authentication methods:
  - **Browser challenge:** Uses Kerberos or NT LAN Manager (NTLM)
  - **Web form:** Uses multi-factor authentication (MFA), security assertion markup language (SAML) single sign-on (SSO), Kerberos, terminal access controller access control system plus (TACACS+), remote authentication dial-in user service (RADIUS), LDAP, or local authentications
  - Client CA
- **GlobalProtect:** Mobile users have an application running on their endpoint for which they must enter login credentials for VPN access to the firewall. The login information is used for User-ID mapping. GlobalProtect is the most recommended method to map device IP addresses to usernames.
- **XML API:** The PAN-OS XML API is used in cases where standard user mapping methods might not work—for example, as third-party VPNs or 802.1x-enabled wireless networks.
- **Client probing:** Client probing is used in a Microsoft Windows environment where the User-ID agent probes client systems by using Windows Management Instrumentation or NetBIOS. Client probing is not a recommended method for user mapping.

The following figure shows the main functionality of the User-ID agent.



The PAN-OS software can use multiple information sources simultaneously to keep an accurate and up-to-date table of the IP-to-user mappings for active sessions.

#### 1.4.2 Determine if User-ID agent or agentless should be used

##### Use Agentless (PAN-OS)

- If you have a small-to-medium deployment with few users and 10 or fewer domain controllers or exchange servers
- If you want to share the PAN-OS-sourced mappings from Microsoft Active Directory (AD), Captive Portal, or GlobalProtect with other PA devices (maximum 255 devices)

##### Use User-ID Agent (Windows)

- If you have a medium-to-large deployment with many users or more than 10 domain controllers
- If you have a multi-domain setup with a large number of servers to monitor

#### 1.4.3 Compare and contrast User-ID agents

##### Identifying the User-ID Agent to Deploy

User-ID has two agents that can be used to monitor servers and gather User-ID information: a built-in agent inside the PAN-OS firewall and a Windows-based client. The built-in agent is called an integrated agent. The Windows-based client can be installed on Windows Server 2008 or later systems. Both agents have the same functionality. Several factors can help determine which agent to use.

An organization might choose to use the Windows agent if it has more than 100 domain controllers, because neither type of agent can monitor more than 100 domain controllers or 50 syslog servers.

Another reason to choose the Windows agent over the integrated PAN-OS agent is to save processing cycles on the firewall's management plane.

However, if network bandwidth is an issue, you might want to use the PAN-OS integrated agent. The PAN-OS integrated agent communicates directly with the servers, while the Windows agent communicates with the servers and then communicates User-ID information to the firewall so that it can update the firewall database.

#### 1.4.4 Methods of User-ID redistribution

##### **Methods of User-ID Redistribution**

Every firewall that enforces a user-based policy requires user mapping information. In a large-scale network, instead of configuring all the firewalls to directly query the mapping information sources, you can streamline resource usage by configuring some firewalls to collect mapping information through redistribution. Redistribution also enables the firewalls to enforce user-based policies when users rely on local sources for authentication (such as regional directory services) but need access to remote services and applications (such as global data center applications). The Data Redistribution feature allows a firewall to be a source of IP user mappings, among other types of data, for any device that is configured to communicate with the agent service of that source firewall or via Panorama.

If you configure an Authentication policy, your firewalls must also redistribute the authentication timestamps that are generated when users are authenticated to access applications and services. Firewalls use the timestamps to evaluate the timeouts for the Authentication policy rules. The timeouts allow a user who successfully authenticates to later request services and applications without authenticating again within the timeout periods. The redistribution of timestamps enables you to enforce consistent timeouts across all the firewalls in your network.

Firewalls share user mappings and authentication timestamps as part of the same redistribution flow; you do not have to configure redistribution separately for each information type.

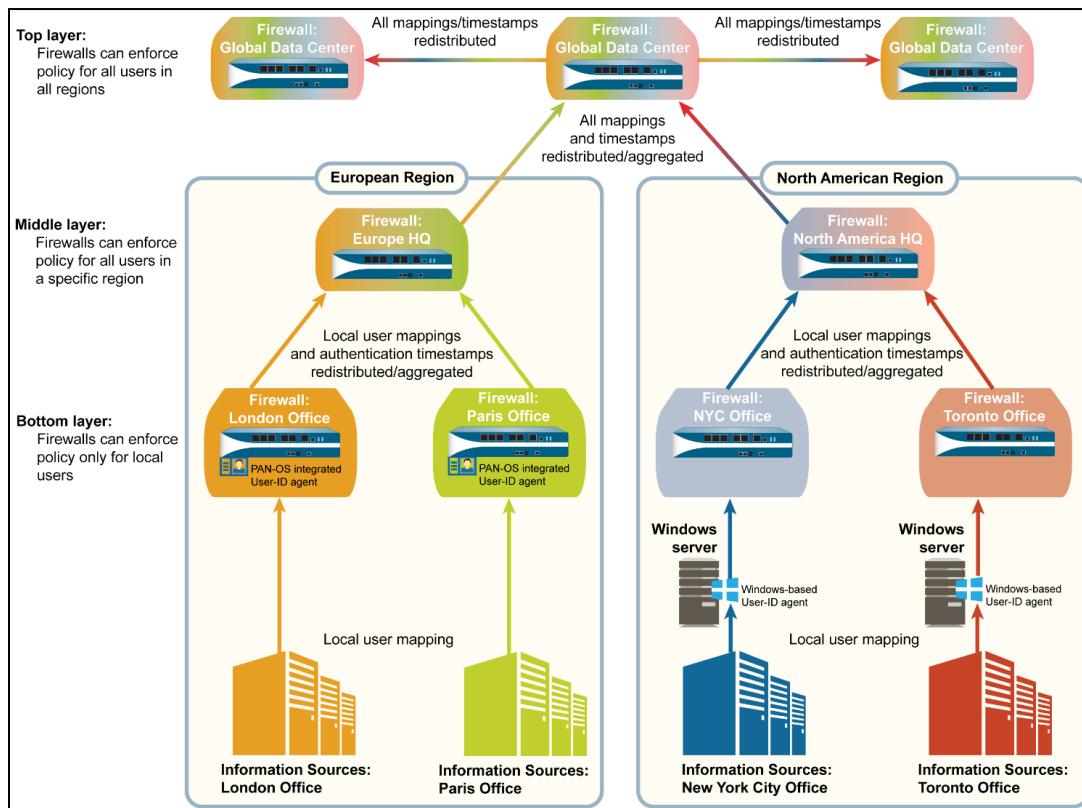
##### **User-ID Table Sharing**

You can enable a firewall or virtual system to serve as a data distribution agent that redistributes user mapping information along with the timestamps associated with authentication challenges. Simply configure the Data Redistribution settings to create an agent that can communicate with any firewalls or other devices to share local information.

##### **User-ID Table Consumption**

To map IP addresses to usernames, User-ID agents monitor sources such as directory servers. The agents send the user mappings to firewalls, Log Collectors, or Panorama. Each appliance can then serve as a redistribution point that forwards the mappings to other firewalls, Log Collectors, or Panorama. Before a firewall or Panorama can collect user mappings, you must configure its connections to the User-ID agents or redistribution points.

## Use Case Example



### 1.4.5 Methods of group mapping

The following are the best practices for group mapping in an AD environment:

- If you have a single domain, you need only one group mapping configuration with an LDAP server profile that connects the firewall to the domain controller with the best connectivity. You can add up to four domain controllers within an LDAP server profile for redundancy. Note that you cannot increase redundancy beyond four domain controllers for a single domain by adding multiple group mapping configurations for that domain.
- If you have multiple domains or multiple forests, you must create a group mapping configuration with an LDAP server profile that connects the firewall to a domain server in each domain or forest. Take steps to ensure that you have unique usernames in separate forests.
  - If you have universal groups, create an LDAP server profile to connect to the root domain of the global catalog server on port 3268 or 3269 for SSL. Then, create another LDAP server profile to connect to the root domain controllers on port 389 or 636 for SSL. This helps ensure that both user and group information is available for all the domains and subdomains.
  - Before using group mapping, configure a primary username for user-based security policies because this attribute identifies users in the policy configuration, logs, and reports.

## 1.4.6 Server profile and authentication profile

### Multi-Factor Authentication Server Profile

A multi-factor authentication (MFA) server profile defines the access method, location, and authentication for integrated MFA vendors. The **MFA Vendor** drop-down list shows supported vendors. A Certificate Profile is required to support the certificate that is used to validate the certificate used by the MFA solution to secure its communication with the firewall.

The screenshot shows the 'Multi Factor Authentication Server Profile' configuration window. It includes fields for 'Profile Name' and 'Certificate Profile'. Under 'Server Settings', there is an 'MFA Vendor' dropdown menu currently set to 'Duo v2'. Below this is a table with the following data:

NAME	
API Host	Okta Adaptive
Integration Key	PingID
Secret Key	RSA SecurID Access
Timeout (sec)	30 [5 - 600]
Base URI	/auth/v2

At the bottom of the window are 'OK' and 'Cancel' buttons.

### Authentication Profile

An Authentication Profile specifies the authentication type and Server Profile for the first Captive Portal-driven authentication. The **Factors** tab incorporates the integrated MFA vendor defined in the multi-factor authentication server profile. Multiple factors can be added that require the user to pass each challenge from the top down.

The screenshot shows the 'Authentication Profile' configuration window, specifically the 'Factors' tab. It includes a 'Profile Name' field and tabs for 'Authentication', 'Factors', and 'Advanced'. A checkbox labeled 'Enable Additional Authentication Factors' is checked, with a note: 'The factors below are used only for Authentication Policy'. Below this is a list area for factors, currently empty, with a header 'FACTORS' and a collapse icon. At the bottom of the list area are controls: '+ Add', '- Delete', '↑ Move Up', and '↓ Move Down'. 'OK' and 'Cancel' buttons are at the bottom right.

## 1.4.7 References

- Map Users to Groups,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/user-id/map-users-to-groups>
- Device > User Identification > Group Mapping Settings,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/user-identification/device-user-identification-group-mapping-settings>
- Group Mapping,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/user-id/user-id-concepts/group-mapping>
- User-ID,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/user-id>
- Best Practices for Securing User-ID Deployments,  
<https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIVPCA0>
- Redistribute Data and Authentication Timestamps,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/user-id/deploy-user-id-in-a-large-scale-network/redistribute-user-mappings-and-authentication-timestamps>
- Data Redistribution Using Panorama,  
<https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/panorama-overview/user-id-redistribution-using-panorama>
- Redistribution,  
<https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-web-interface-help/user-identification/device-user-identification-user-mapping/user-id-agent-setup/user-id-agent-setup-redistribution>
- Device > Data Redistribution,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/device/device-data-redistribution>

## 1.5 Determine how and when to use the Authentication policy

### 1.5.1 Purpose of, and use case for, the Authentication policy

#### **MFA and Authentication Policy**

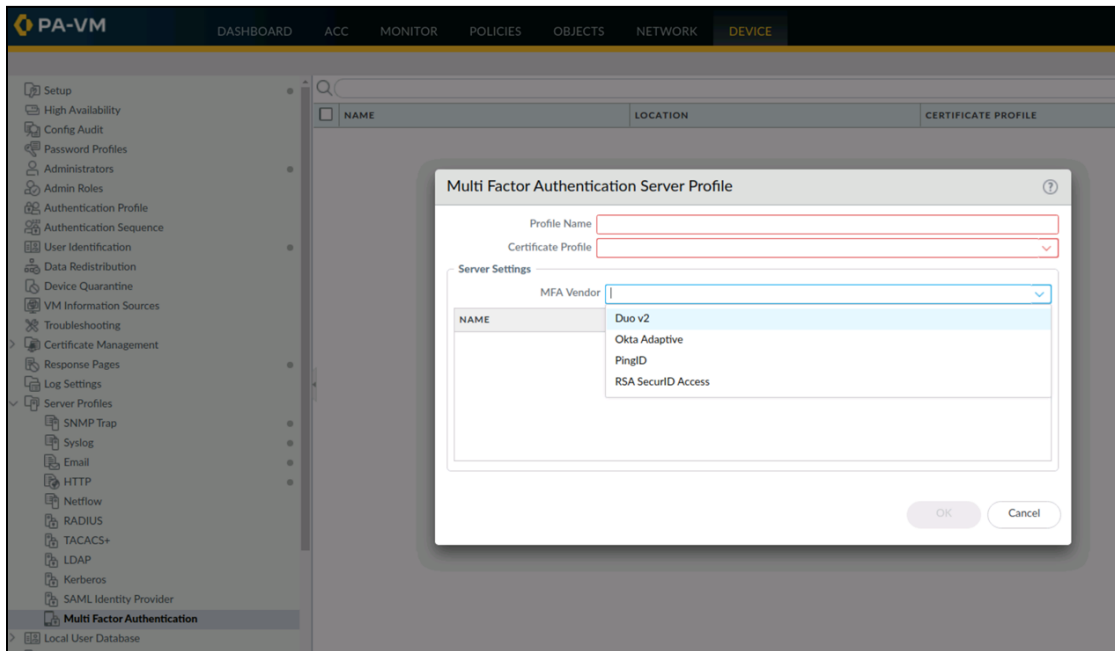
You can configure MFA to ensure that each user authenticates using multiple methods (factors) when they access highly sensitive services and applications. For example, you can force users to enter a login password and then a verification code that they receive by phone before allowing access to key financial documents. This approach helps to prevent attackers from accessing every service and application in your network just by stealing passwords.

For end-user authentication via the Authentication policy, the firewall directly integrates with several MFA platforms (such as Duo v2, Okta Adaptive, PingID, and RSA SecurID) and integrates through RADIUS with other MFA platforms.

MFA is driven by an Authentication policy that allows the precise application of appropriate authentication. These policy rules can invoke simple Captive Portal challenge pages for one-time

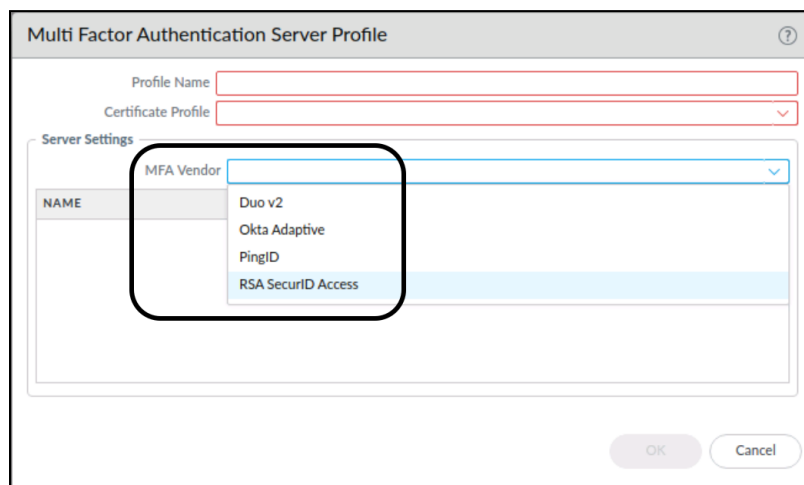
authentication or can include one (or more) integrated MFA vendor Server Profiles that are included in the Authentication Profiles for additional challenges.

After a user successfully completes all of the challenges, Security policy rules that allow access to that protected service are evaluated.



### Special Note About MFA

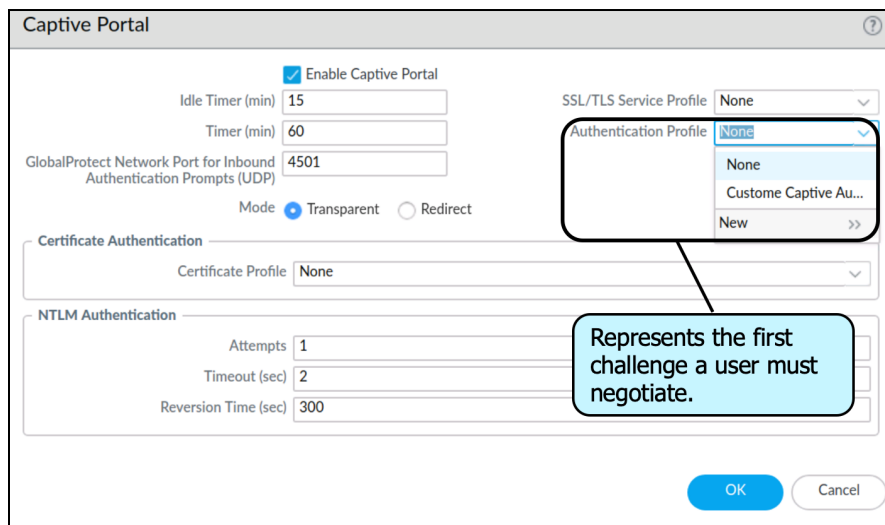
The Palo Alto Networks firewalls support MFA. A multi-factor authentication server profile is used to natively integrate a firewall with an external third-party MFA solution. The MFA factors that the firewall supports include push, Short Message Service (SMS), voice, and one-time password (OTP) authentication. This profile identifies the specific product with its configuration information.



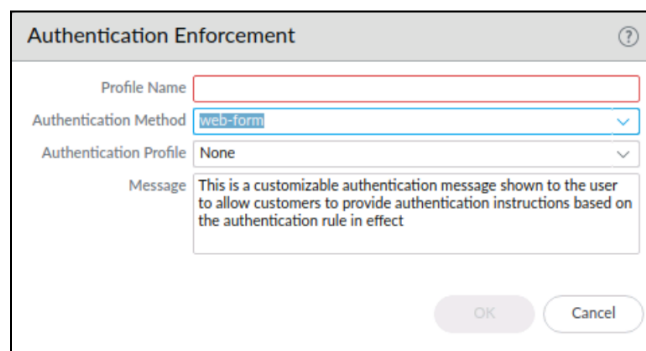


The Multi Factor Authentication Server Profile can be a part of multiple authentication challenges that a user must respond to. For example, you can force users to enter a login password and then enter a verification code that they receive by phone before they can access critical financial documents.

The firewall challenges a user with a Captive Portal. Captive Portal configuration includes an Authentication Profile selected for base configuration that represents the first challenge that a user must negotiate.



An Authentication Enforcement policy is then used to join the MFA product as a second required authentication. The selection of the MFA product's Authentication Profile adds it as a second authentication requirement for users.



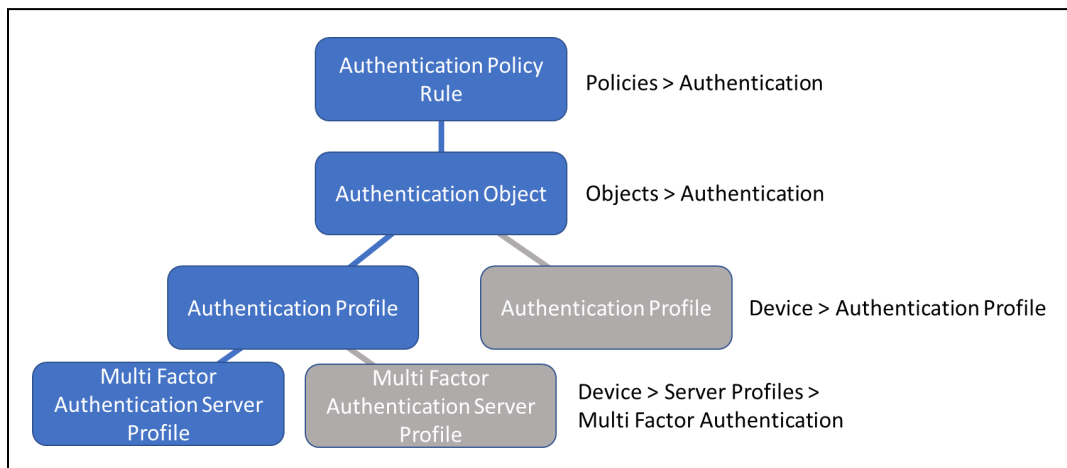
## 1.5.2 Dependencies

### Dependencies for Implementing MFA

Before you use MFA to protect sensitive services and applications, you must configure several settings in the Palo Alto Networks firewall. MFA authentication is triggered when a user requests access to a service that appears in the traffic that the firewall processes. The traffic is first evaluated by an Authentication policy rule. When a match is found, the authentication action for the rule is taken.

The screenshot shows the 'Authentication Policy Rule' configuration window. It has tabs for 'General', 'Source', 'Destination', 'Service/URL Category', and 'Actions'. The 'Actions' tab is selected. The 'Authentication Enforcement' dropdown is set to 'Enforcement'. The 'Timeout (min)' is set to '60'. Under 'Log Settings', the 'Log Authentication Timeouts' checkbox is unchecked, and 'Log Forwarding' is set to 'None'. 'OK' and 'Cancel' buttons are at the bottom right.

The following figure shows the relationship of the required objects to configure the Authentication policy rule.



- **Authentication Enforcement object:** This object specifies which Authentication Profile to use and is assigned to an Authentication policy rule. A Captive Portal authentication method is also specified. A custom message can be included for the user explaining how to respond to the challenge.

### 1.5.3 Captive portal versus GlobalProtect (GP) client

#### Captive Portal

If the firewall or the User-ID agent can't map an IP address to a username—for example, if the user isn't logged in or uses an operating system such as Linux that your domain servers don't support—you can configure Captive Portal. Any web traffic (HTTP or HTTPS) that matches a Captive Portal policy rule requires user authentication. You can base the authentication on a transparent browser-challenge (Kerberos SSO or NTLM in Captive Portal authentication), web form (for RADIUS, TACACS+, LDAP, Kerberos, or local database authentication), or client certificates.

#### GlobalProtect client

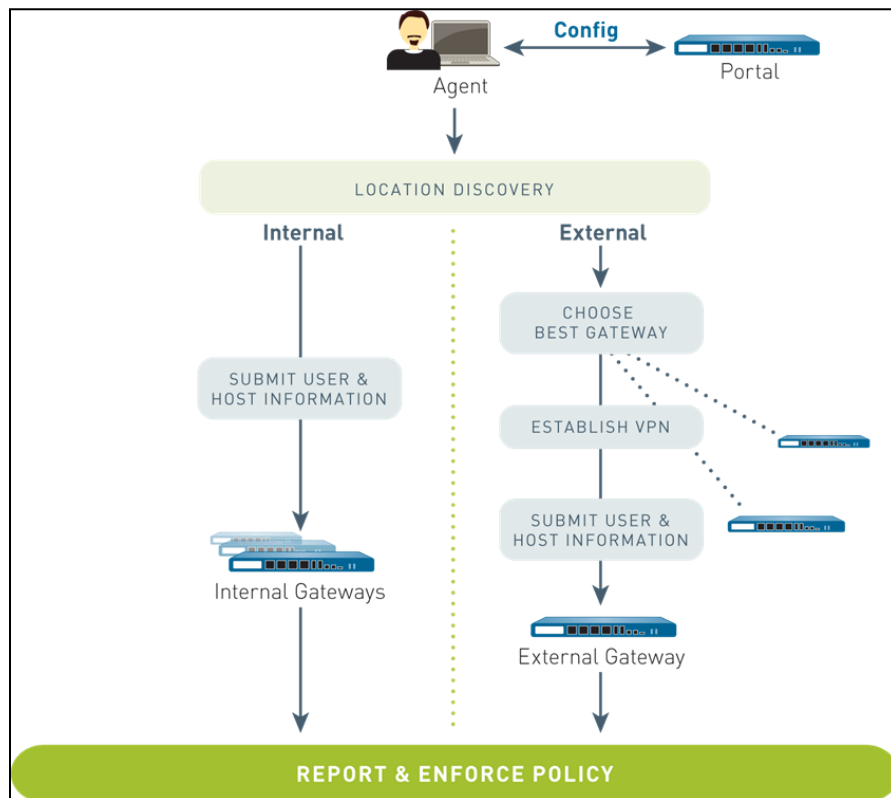
The GlobalProtect client software runs on end user systems. It shares User-ID information and enables access to network resources via the GlobalProtect portals and gateways you have deployed.

#### The GlobalProtect App

The GlobalProtect app runs on Windows, macOS, iOS, Android, Linux, Chromebook, and IoT Devices.

On Microsoft Windows and Apple macOS devices, the application can be customized by changing behavior and hiding certain UI elements to suit a particular environment's requirements.

The following diagram illustrates how the GlobalProtect portals, gateways, and agents/apps work together to enable secure access for all the users, regardless of the devices they are using or their location.



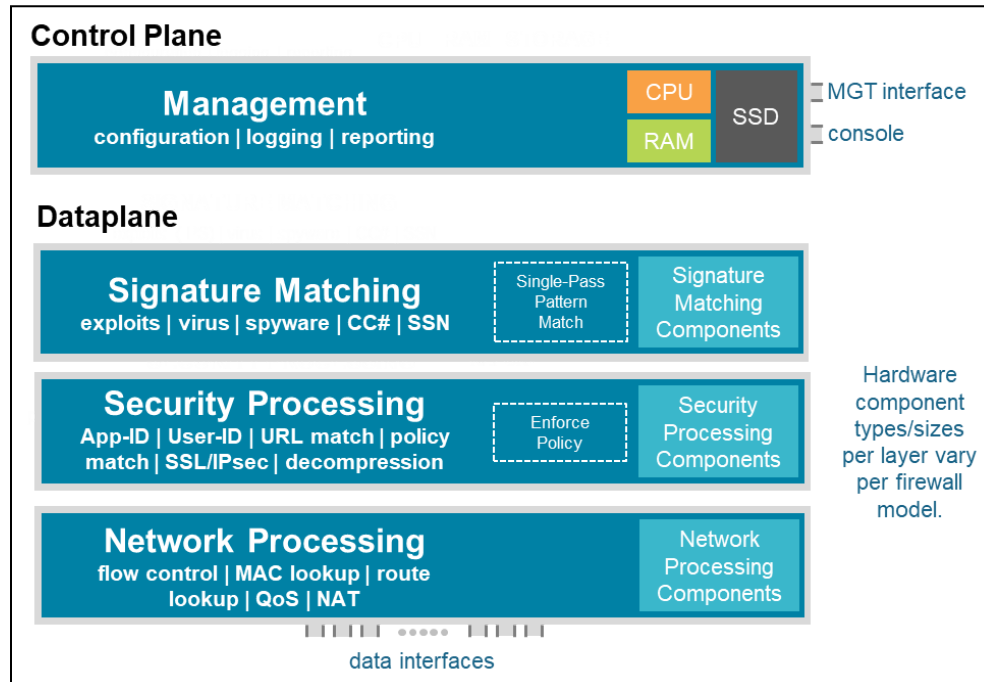
## 1.5.4 References

- Configure Authentication Portal, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/user-id/map-ip-addresses-to-users/map-ip-addresses-to-username-using-captive-portal/configure-captive-portal>
- Configure Multi-Factor Authentication, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/authentication/configure-multi-factor-authentication>
- Map IP Addresses to Usernames Using Authentication Portal, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/user-id/map-ip-addresses-to-users/map-ip-addresses-to-username-using-captive-portal>
- Authentication Policy, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/authentication/authentication-policy>

## 1.6 Differentiate between the fundamental functions that reside on the management plane and data plane

### Management Planes and Data Planes

Whether the management plane and data plane functionality is physical or virtual, it is integral to all the Palo Alto Networks firewalls. On physical firewalls, these functions have dedicated hardware resources, which makes them independent of each other. On virtual firewalls, these are still logically segregated. The following figure details the architecture of a PA-220 firewall:



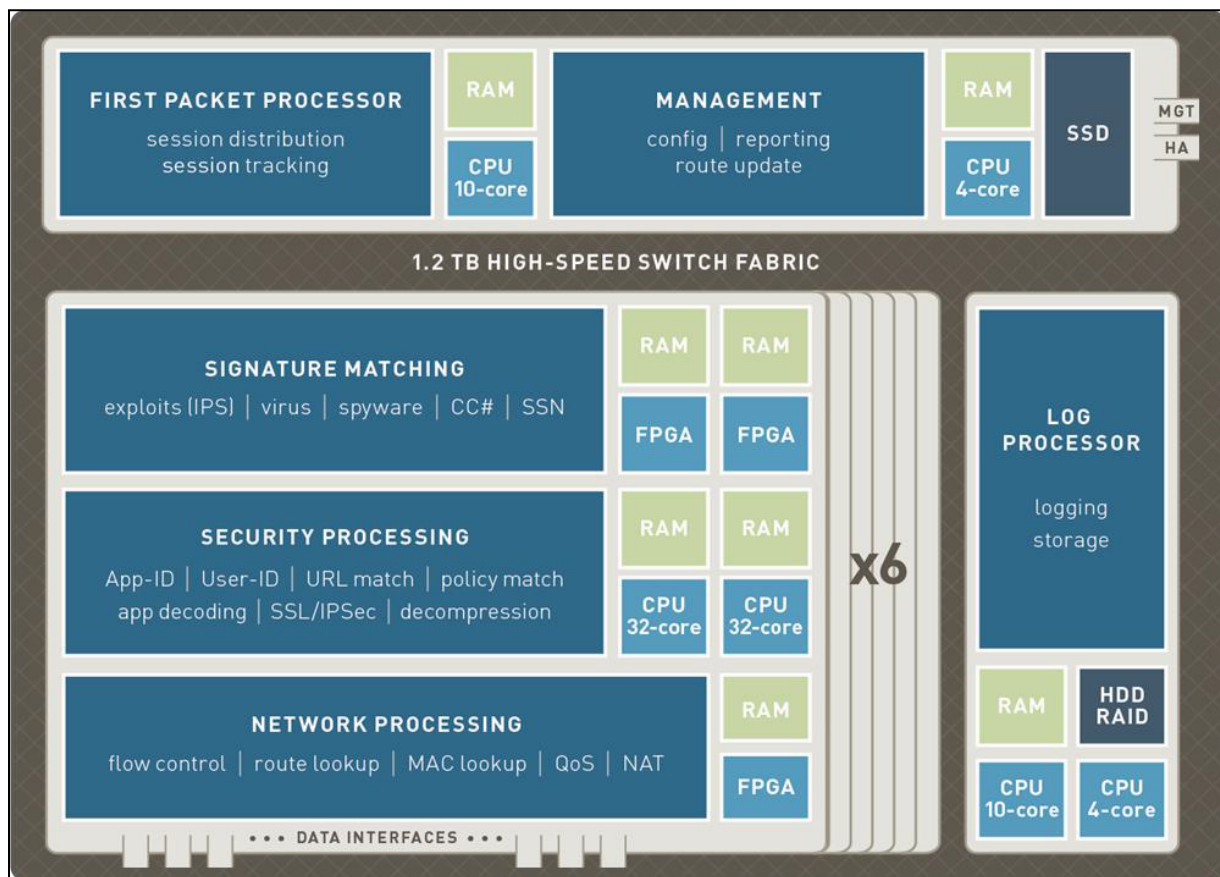
Palo Alto Networks maintains the management plane and data plane separation to protect system resources.

Every Palo Alto Networks firewall assigns a minimum of these functions to the management plane:

- Configuration management
- Logging
- Reporting functions
- User-ID agent process
- Route updates

The management network terminates directly on this plane, as does the console connector on physical firewalls.

On the PA-7000 Series firewalls, dedicated log collection and processing is implemented on a separate card. The following figure provides an overview of the PA-7000 Series architecture:



### Identify the functions that reside on the data plane

The following functions are assigned to the data plane:

- Signature match processor
- All Content-ID and App-ID services
- Security processors
- Session management

- Encryption and decryption
- Compression and decompression
- Policy enforcement
- Network processor
- Route
- Address Resolution Protocol (ARP)
- MAC lookup
- QoS
- NAT
- Flow control

The data plane connects directly to the traffic interfaces. As more computing capability is added to more powerful firewall models, the management planes and data planes gain other functionality, as required, sometimes implemented on dedicated cards. Several core functions gain field-programmable gate arrays (FPGAs) or custom application-specific integrated circuits for flexible high-performance processing. Additional management plane functions might include the following:

- First packet processing
- Switch fabric management

### **Scope the impact of using SSL decryption**

When decryption is performed correctly, it enhances security and prevents adversaries from misusing encrypted traffic to attack an organization. Decrypting traffic can weaken security, but if you follow best practices, decryption provides visualization requirements into all of the traffic, improving the security posture. Decryption also protects from adversaries that hide threats in encrypted tunnels.

### **Scope the impact of turning logs on for every Security policy**

Visibility is important, and having logs enabled for all Security policies assists with this. By default, traffic that hits the default policies do not get logged. To change this behavior, you need to overwrite the intrazone and interzone default rules.

Enabling logging on the default rules can generate a large amount of logs and should only be done if required and if the potential impact to storage capacity and performance has been adequately analyzed.

## **1.6.1 References**

- SSL Decryption Series: The Security Impact of HTTPS Interception, <https://blog.paloaltonetworks.com/2018/10/ssl-decryption-series-security-impact-https-interception/>
- Size the Decryption Firewall Deployment, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/prepare-to-deploy-decryption/size-the-decryption-firewall-deployment>
- Tune or Reduce Firewall Logs, <https://splunk.paloaltonetworks.com/tune-or-reduce-firewall-logs.htm>

## 1.7 Define multiple virtual systems (multi-vsys) environment

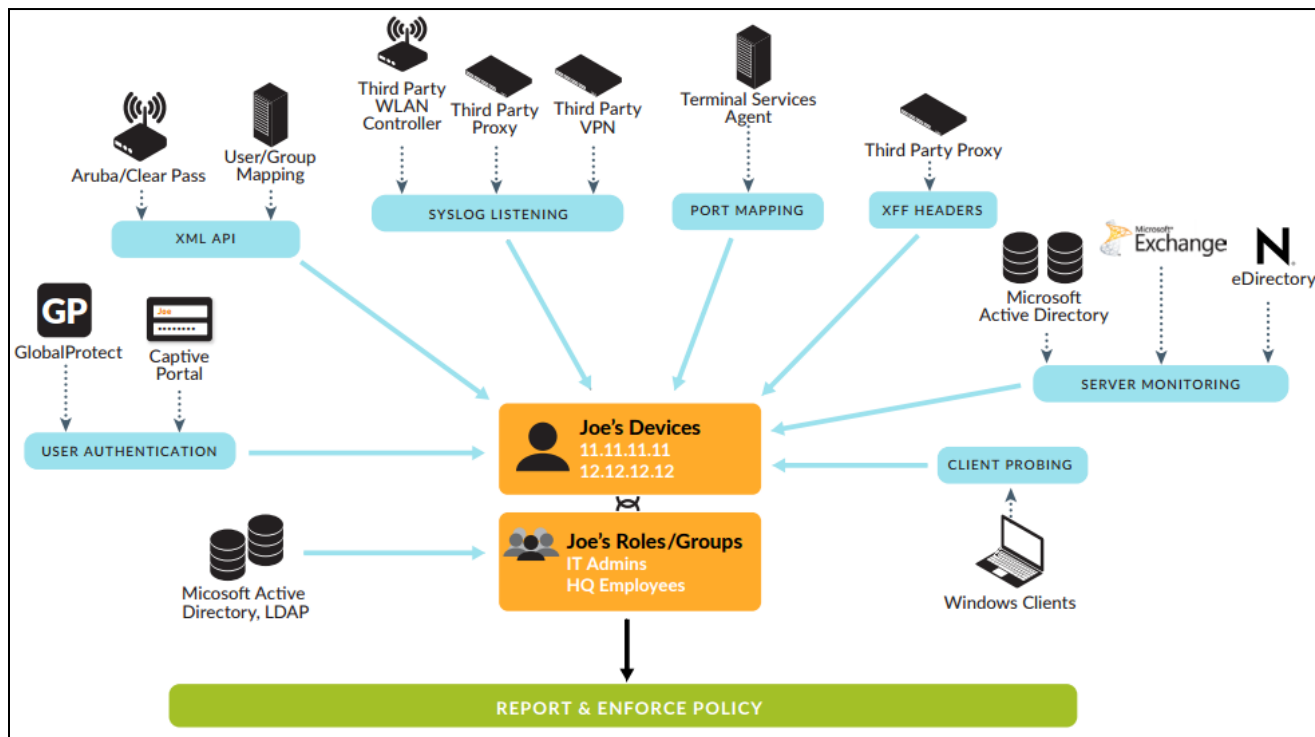
### 1.7.1 User-ID hub

User-ID enables you to identify all the users on your network by using a variety of techniques to ensure that you can identify users across all of the locations by using a variety of access methods and operating systems, including Microsoft Windows, Apple iOS, Mac OS, Android, and Linux/UNIX. Knowing who your users are instead of just their IP addresses enables:

- **Visibility** — Improved visibility into application usage based on users gives you a more relevant picture of network activity. The power of User-ID becomes evident when you notice a strange or unfamiliar application on your network. Using either ACC or the Log Viewer, your security team can identify the application, the user, the bandwidth and session consumption, the source and destination of the application traffic, and any associated threats.
- **Policy control** — Tying user information to Security policy rules improves the safe enablement of applications traversing the network and ensures that only those users who have a business need for an application have access. For example, some applications, such as the SaaS applications that enable access to Human Resources services (such as Workday or ServiceNow) must be available to any known user on the network. However, for more sensitive applications, you can reduce the attack surface by ensuring that only users who need these applications can access them. For example, while IT support personnel may legitimately need access to remote desktop applications, the majority of users do not.
- **Logging, reporting, forensics** — If a security incident occurs, forensics analysis and reporting based on user information rather than just IP addresses provide a more complete picture of the incident. For example, you can use the predefined User/Group Activity to see a summary of the web activity of individual users or user groups, or view the SaaS Application Usage report to see which users are transferring the most data over unsanctioned SaaS applications.

To enforce user- and group-based policies, the firewall must be able to map the IP addresses, in the packets it receives, to the usernames. User-ID provides many mechanisms to collect this User Mapping information. For example, the User-ID agent monitors server logs for login events and listens for syslog messages from authenticating services. To identify mappings for IP addresses that the agent didn't map, you can configure the Authentication Policy to redirect HTTP requests to a Captive Portal login. You can tailor the user mapping mechanisms to suit your environment and even use different mechanisms at different sites to ensure that you enable safe access to the applications for all the users, in all the locations, all the time.

## User-ID



To enable user- and group-based policy enforcement, the firewall requires a list of all the available users and their corresponding group memberships so that you can select groups when defining your policy rules. The firewall collects Group Mapping information by connecting directly to the LDAP directory server or by using XML API integration with the directory server.

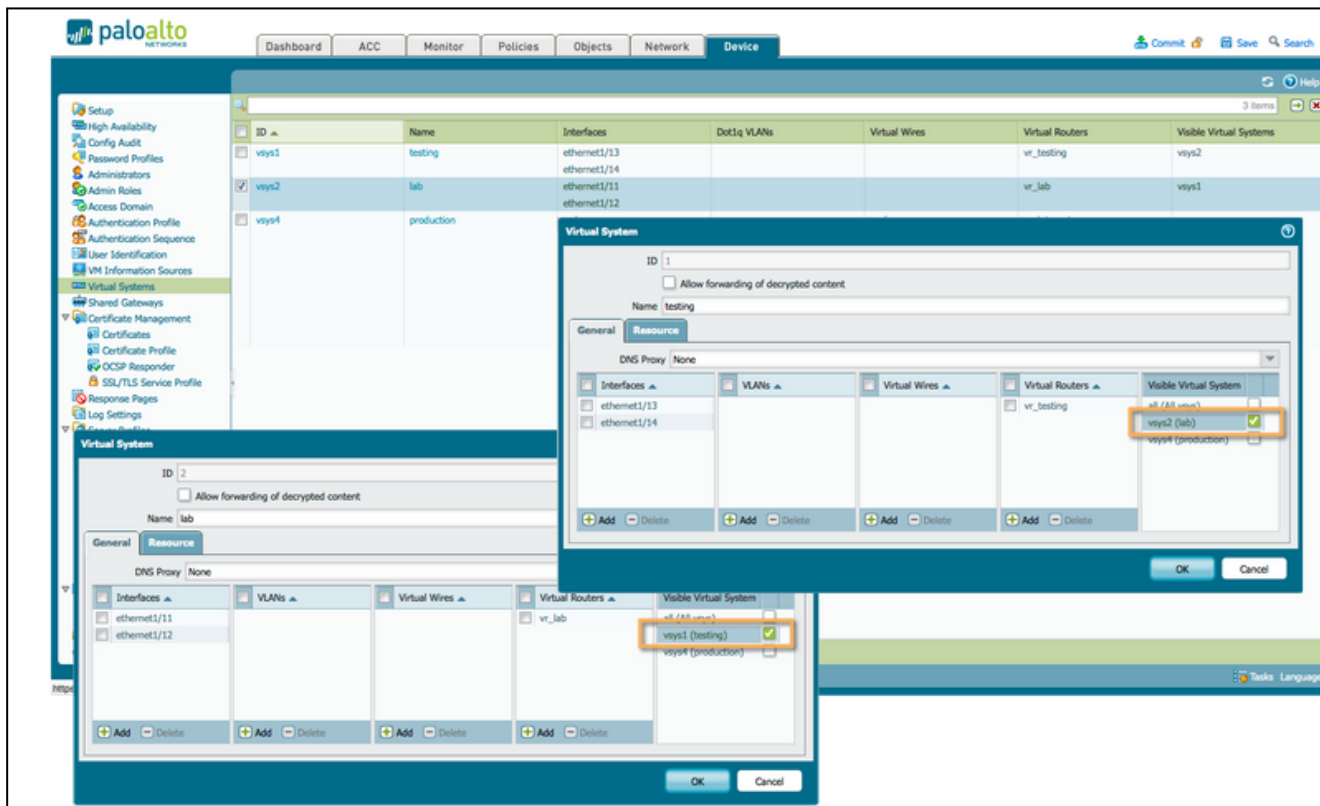
### 1.7.2 Inter-vsyst routing

A virtual system (vsys) is a separate, logical firewall instance within a single physical chassis. Enabling virtual systems on the firewall can help logically separate physical networks from each other. Separate networks are handy when specific networks should not be connected to each other. Using virtual systems also allows you to control which administrators can control certain parts of the network and firewall configuration.


In some cases, however, some connectivity needs to be enabled between virtual systems. Rather than physically connecting separate networks, which could cause a potential security breach, limited routing can be enabled to allow only specific subnets to communicate. The Security policy can then be applied to prevent abuse of this bridge between networks.



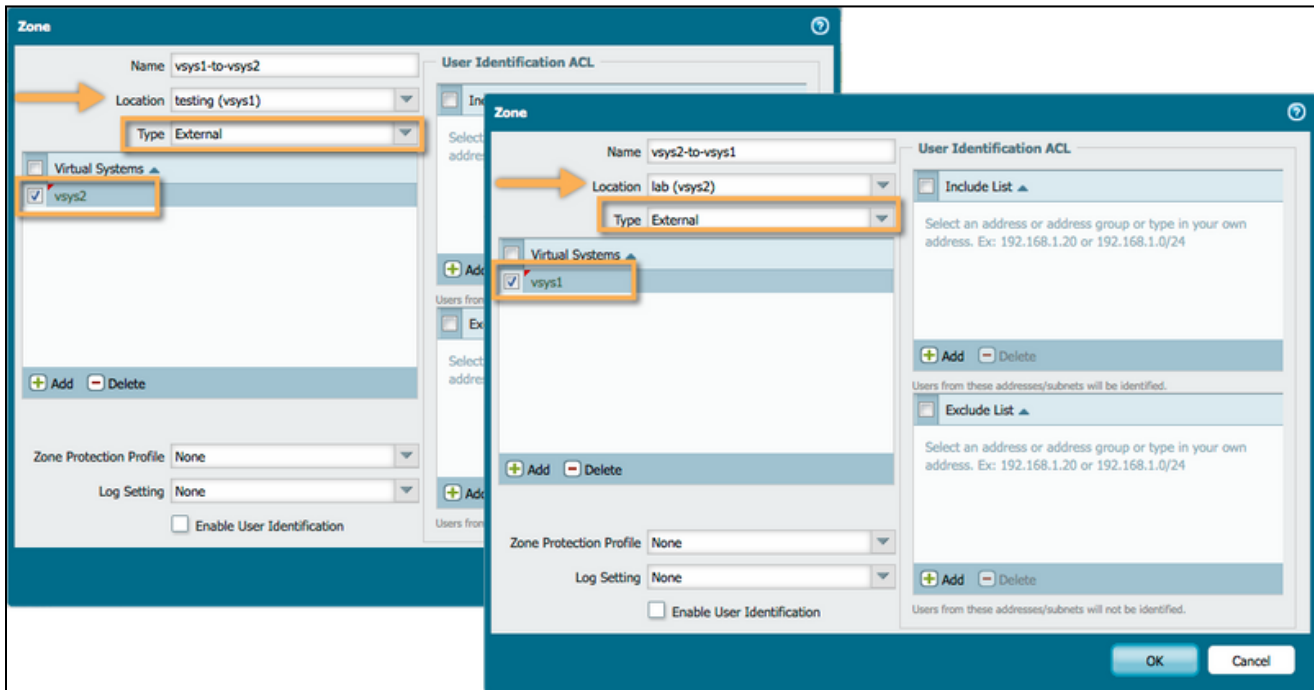
First, visibility should be enabled between virtual systems. Since a vsys acts as a standalone system, it is not aware of any other vsys residing on the same physical chassis.



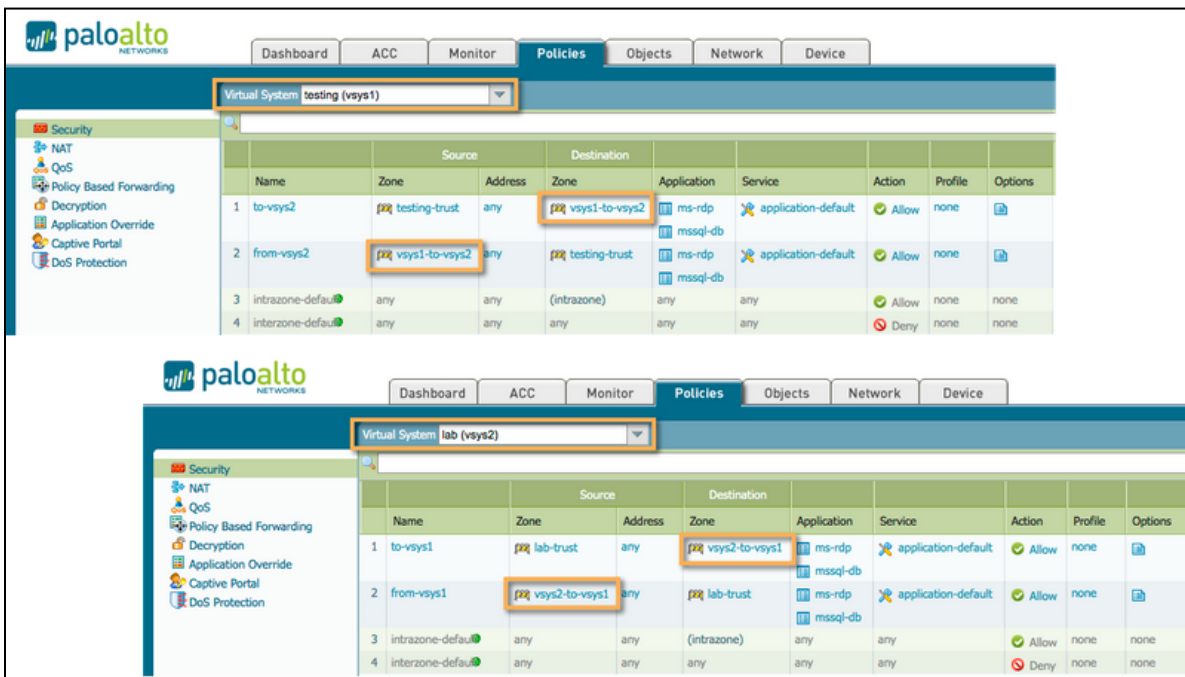
Next, a new type of zone called “External” should be created on each vsys to allow sessions to traverse into a zone that connects the virtual systems. The **External** type forms a network of sorts that allows the virtual systems to communicate.



On each participating vsys, create a zone with type **External**. Add the destination Virtual System to allow this zone to represent the remote vsys. Multiple destination virtual systems can be added.



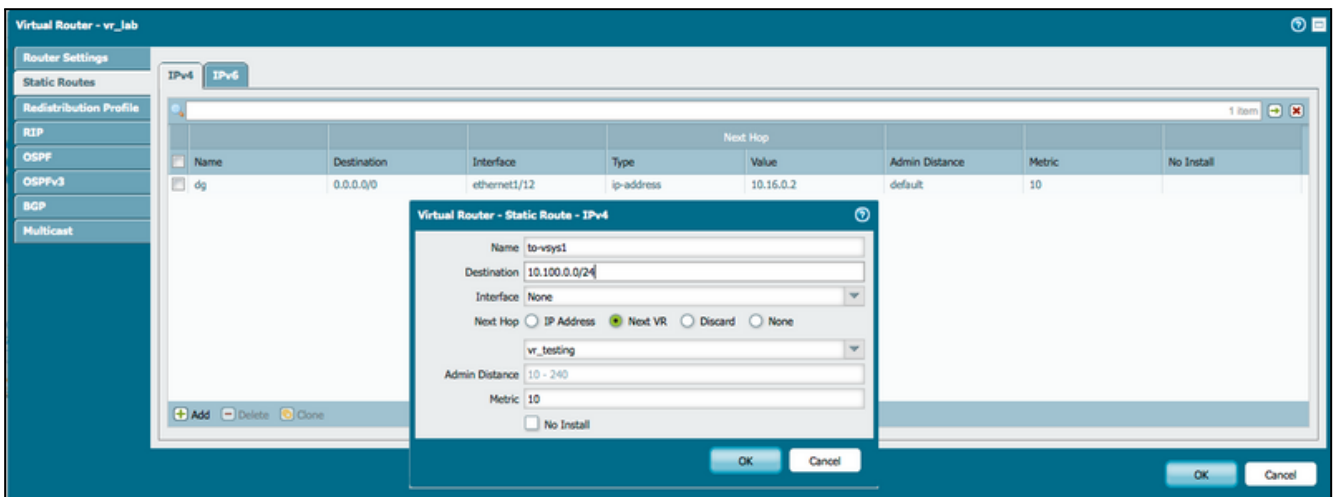
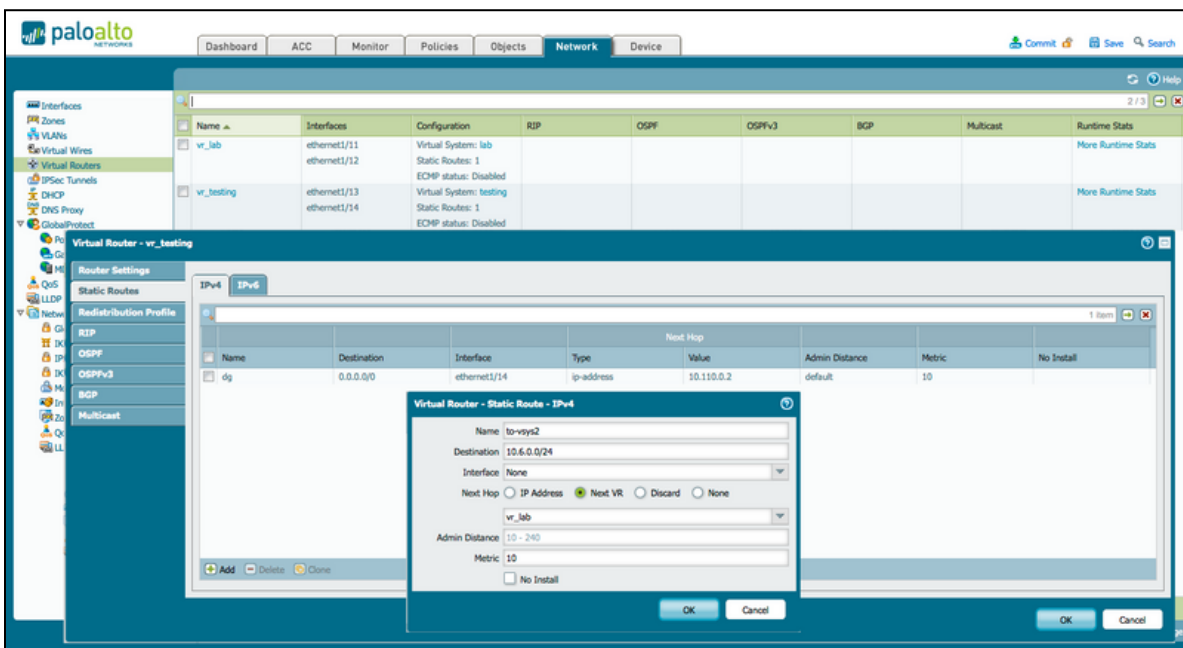
Each vsys should then be configured with a Security policy that allows the local zone to connect out to the External zone or from the External zone to the trusted network if the connection is to be considered inbound. For example, in the case of an out-of-the-box (OOB) network, the IT-vsys can be allowed an outbound connection to the External zone and the OOB vsys could allow an inbound connection from the External zone.



Since the virtual routers are not aware of the subnets available in the remote virtual systems, routing needs to be added to direct traffic properly to the External zone.

ethernet1/11	Layer3		10.6.0.1/24	vr_lab	Untagged	none	lab	lab-trust
ethernet1/12	Layer3		10.16.0.1/24	vr_lab	Untagged	none	lab	lab-untrust
ethernet1/13	Layer3		10.100.0.1/24	vr_testing	Untagged	none	testing	testing-trust
ethernet1/14	Layer3		10.110.0.1/24	vr_testing	Untagged	none	testing	testing-untrust

Configure each Virtual Router to be configured with routes for the appropriate remote subnets, with the next hop set to the remote virtual router of the virtual system. In the example, the 'testing' virtual router needs to be configured with a static route for the lab-trust subnet 10.6.0.0/24 pointing to the vr\_lab virtual router, and a return route on the vr\_lab virtual router for testing-trust subnet 10.100.0.0/24 pointing to the vr\_testing remote virtual router.



---

When this configuration is committed, the clients located in the trust zones of both vsys1 and vsys2 are able to connect to each other through the Microsoft Remote Desktop or MSSQL applications, as per the security policy.

### 1.7.3 Service routes

The firewall uses the management (MGT) interface, by default, to access external services, such as DNS servers, external authentication servers, and Palo Alto Networks services such as software, URL updates, licenses, and AutoFocus. An alternative to using the MGT interface is to configure a data port (a regular interface) to access these services. The path from the interface to the service on a server is known as a service route. The service packets exit the firewall on the port assigned for the external service, and the server sends its response to the configured source interface and source IP address.

You can [configure service routes](#) globally for the firewall or [customize service routes for a virtual system](#) on a firewall enabled for multi-vsys so that you have the flexibility to use interfaces associated with a vsys. Any vsys that does not have a service route configured for a particular service inherits the interface and IP address that are set globally for that service.

### 1.7.4 References

- User-ID Overview,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/user-id/user-id-overview>
- Inter VSYS Routing,  
<https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CISVCA0>
- Service Routes Overview,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/service-routes/service-routes-overview>

## Domain 2: Deploy and Configure Core Components

### 2.1 Configure Management Profiles

Use Interface Management profiles to configure in-band interfaces for allowing management access and other features. By default, for the benefit of a stronger security stance, dataplane interfaces do not allow any administration to terminate on them. This includes HTTP/HTTPS, Ping, SNMP, and others. To enable these features, create an Interface Management profile with the appropriate features enabled and assign the profile to the required interface.

#### 2.1.1 Interface Management Profile

To configure an Interface Management Profile, perform the following steps:

1. Navigate to **Network > Network Profiles > Interface Mgmt** and click **Add**.
2. Select the network protocols that the interface permits (allows) for management traffic. Choose from **Ping, Telnet, SSH, HTTP, HTTP OCSP, HTTPS, or SNMP**.
3. Select the services that the interface permits for management traffic. Choose from the following:
  - a. Response Pages (for Authentication Portal or URL Admin Override)
  - b. User-ID (to redistribute data and authentication timestamps)
  - c. User-ID Syslog Listener-SSL or User-ID Syslog Listener-UDP (to configure User-ID to monitor syslog senders for user mapping over SSL or User Datagram Protocol [UDP] traffic)
4. Optionally, add IP addresses to permit access to the interface. If you don't add an IP address, the interface will have no IP address restrictions.
5. Click **OK**.

#### 2.1.2 SSL/TLS profile

The Palo Alto Networks firewalls and Panorama use the SSL/TLS service profiles to specify a certificate and the allowed protocol versions for SSL/TLS services. The firewall and Panorama use SSL/TLS for the Authentication Portal, the GlobalProtect portals and gateways, the inbound traffic on the MGT interface, the URL Admin Override feature, secure syslog log forwarding, and the User-ID syslog listening service. By defining the protocol versions, you can use a profile to restrict the cipher suites and TLS versions that are available for securing communication with the clients requesting the services. This improves network security by enabling the firewall or Panorama to avoid the SSL/TLS versions that have known weaknesses. If a service request involves a protocol version that is outside the specified range, the firewall or Panorama downgrades or upgrades the connection to a supported version.

#### 2.1.3 References

- Use Interface Management Profiles to Restrict Access, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/configure-interfaces/use-interface-management-profiles-to-restrict-access>
- How to Configure the Management Interface IP, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIN7CAK>

- Network > Network Profiles > Interface Mgmt,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/network/network-network-profiles/network-network-profiles-interface-mgmt>
- Configure an SSL/TLS Service Profile,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/certificate-management/configure-an-ssl-tls-service-profile>

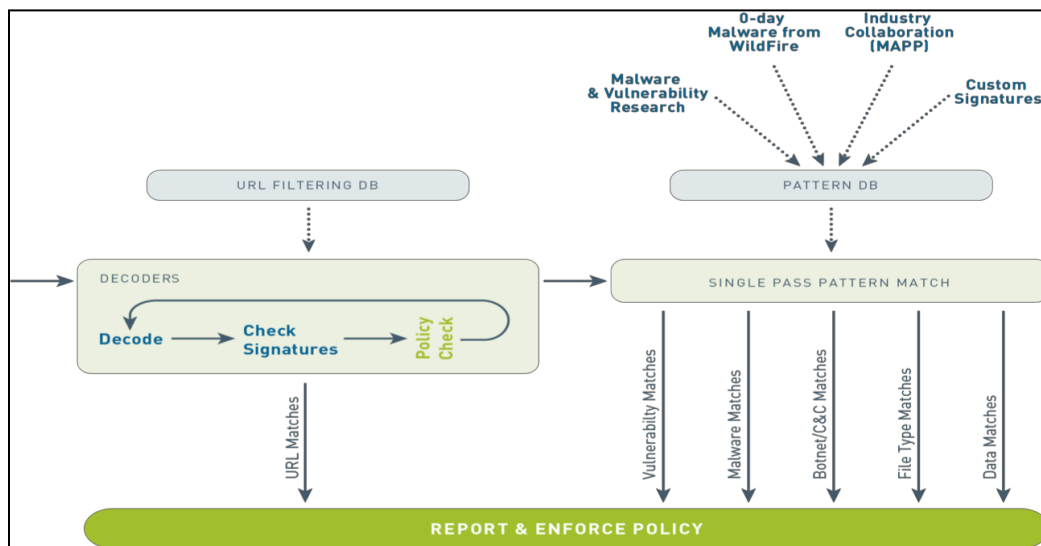
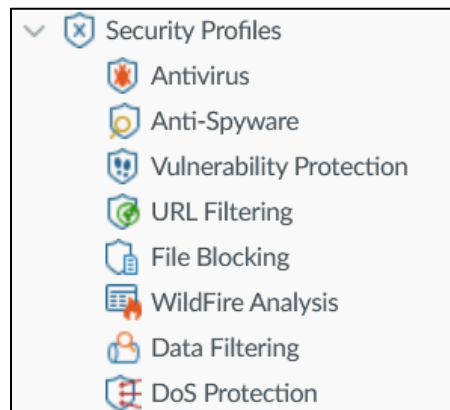
## 2.2 Deploy and configure Security Profiles


Security Profiles implement specific protections provided by the Palo Alto Networks Content-ID next-generation technology. After Security Profiles are created, they are attached to Security policy rules specifying the Content-ID scans to be performed on the traffic allowed by a policy rule. These profiles must be attached to the Security policy rules to invoke their protections and will be applied only to the traffic handled by that particular rule.

### 2.2.1 Custom configuration of different Security Profiles and Security Profile Groups

#### Security Profile Overview

Security Profiles include:





All scanning is done by signature matching on a streaming basis (not on a file basis). These signatures are updated based on the configuration and licensing options. For example, with a WildFire license, new virus and malware signatures can be installed in real time. If the firewall has a Threat Prevention license but no WildFire license, signatures from WildFire would be updated only every 24 hours.

Content scanning consumes firewall resources after it is enabled. Consult a firewall comparison chart to identify the model with appropriate “Threat Enabled” throughput.

### **Identifying Security Profiles for use**

Although Security policy rules enable you to allow or block traffic on your network, Security Profiles help you define an allow-but-scan rule, which scans the allowed applications for threats, such as viruses, malware, spyware, and DoS attacks. When traffic matches the allow rule that is defined in the Security policy, the Security Profile(s) attached to the rule are applied for further content inspection, such as antivirus checks and data filtering. Security Profiles are the features that provide the services of the Content-ID feature of PAN-OS software.

Security Profiles are not used in the match criteria of a traffic flow. The Security Profile is applied to scan traffic after the application or category is allowed by the Security policy.

The firewall provides default Security Profiles that you can use out of the box to begin protecting your network from threats. The Security Profiles attached to the Security policy “allow” rules determine the type of threat detection performed on the traffic.

You can add the Security Profiles that are commonly applied together to create a Security Profile Group; this set of profiles can be treated as a unit and added to the Security policy rules in one step (or included in the Security policy rules, by default, if you choose to set up a default Security Profile Group).

### **Antivirus Profiles**

Antivirus Profiles protect against viruses, worms, Trojan horses, and spyware downloads. The Palo Alto Networks antivirus solution uses a stream-based malware prevention engine that inspects traffic the moment the first packet is received to provide protection for clients without significantly impacting firewall performance. This profile scans for a wide variety of malware in executables, PDF files, HTML, and JavaScript viruses, and it includes support for scanning inside compressed files and data-encoding schemes. If you have enabled decryption on the firewall, the profile also enables the scanning of decrypted content.

The default profile inspects all the listed protocol decoders for viruses and generates alerts for the SMTP, IMAP, and POP3 protocols while blocking for FTP, HTTP, and Server Message Block (SMB) protocols. You can configure the action for a decoder or antivirus signature and specify how the firewall responds to a threat event:

- **default:** For each threat signature and antivirus signature that is defined by Palo Alto Networks, a default action is specified internally. Typically, the default action is an alert or a reset-both. The default action is displayed in parenthesis—for example, *default (alert)* in the threat or antivirus signature.
- **allow:** This action permits the application traffic.
- **alert:** This action generates an alert for each application traffic flow. The alert is saved in the Threat log.
- **drop:** This action drops the application traffic.
- **reset-client:** For TCP, this action resets the client-side connection. For UDP, it drops the connection.
- **reset-server:** For TCP, this action resets the server-side connection. For UDP, it drops the connection.
- **reset-both:** For TCP, this action resets the connection on both client and server ends. For UDP, it drops the connection.

Customized profiles can be used to minimize antivirus inspection for traffic between trusted security zones. Customized profiles can also maximize the inspection of traffic received from untrusted zones, such as the internet, along with the traffic sent to highly sensitive destinations, such as server farms.

The Palo Alto Networks WildFire system also provides signatures for persistent threats that are more evasive and have not yet been discovered by other antivirus solutions. As threats are discovered by WildFire, signatures are quickly created and then integrated into the standard antivirus signatures that can be downloaded daily by Threat Prevention subscribers and sub-hourly for WildFire subscribers.

The WildFire inline ML option in the Antivirus profile enables the firewall data plane to apply ML on the PE (portable executable), ELF (executable and linked format), and MS Office files and on the PowerShell and shell scripts in real time. This layer of antivirus protection complements the WildFire-based signatures to provide extended coverage for those files whose signatures do not already exist. Each inline ML model dynamically detects malicious files of a specific type by evaluating file details, including decoder fields and patterns, to formulate a high probability classification of a file. This protection extends to the currently unknown and the future variants of all the threats that match the characteristics that Palo Alto Networks has identified as malicious. To keep up with the latest changes in the threat landscape, inline ML models are added or updated via content releases. Before you can enable WildFire inline ML, you must possess an active WildFire subscription.



## Anti-Spyware Profiles

Anti-Spyware Profiles block spyware on compromised hosts from trying to phone-home or beacon out to the external C2 servers, thus allowing you to detect any malicious traffic leaving the network from infected clients. You can apply various levels of protection between zones. For example, you may want to have custom Anti-Spyware Profiles that minimize inspection between trusted zones while maximizing inspection on the traffic received from an untrusted zone, such as an internet-facing zone.

You can define your own custom Anti-Spyware Profiles or choose one of the following predefined profiles when applying anti-spyware to a Security policy rule:

- **Default:** This profile uses the default action for every signature as specified by Palo Alto Networks when the signature is created.
- **Strict:** This profile overrides the default action of critical-, high-, and medium-severity threats to the block action, regardless of the action defined in the signature file. This profile still uses the default action for low- and informational-severity signatures.

After the firewall detects a threat event, you can configure the following actions in an Anti-Spyware Profile:

- **default**
- **allow**
- **alert**
- **drop**
- **reset-client**
- **reset-server**
- **reset-both**

**Note:** These actions are similar to the firewall response discussed earlier. In some cases, when the profile action is set to reset-both, the associated threat log might display the action as reset-server, which occurs when the firewall detects a threat at the beginning of a session and presents the client with a 503 block page. The block page disallows the connection, so the client side does not need to be reset, only the server-side connection does.

- **Block IP:** This action blocks traffic from either a source or a source-destination pair. It is configurable for a specified period of time.

You can also enable the DNS Sinkholing action in the Anti-Spyware Profiles to enable the firewall to forge a response to a DNS query for a known malicious domain, thus causing the malicious domain name to resolve to an IP address that you define. This feature helps to identify infected hosts on the protected network by using DNS traffic. Infected hosts can then be easily identified in the Traffic and Threat logs because any host that attempts to connect to the sinkhole IP address is most likely infected with malware.

## Vulnerability Protection Profiles

Vulnerability Protection Profiles stop attempts to exploit system flaws or gain unauthorized access to systems. These profiles can highlight when traffic indicates that a server or client is vulnerable. Although Anti-Spyware Profiles help identify infected hosts as traffic leaves the network,

Vulnerability Protection Profiles protect against threats entering the network. For example, Vulnerability Protection Profiles help protect against buffer overflows, illegal code execution, and other attempts to exploit system vulnerabilities. The default Vulnerability Protection Profile protects clients and servers from all known critical-, high-, and medium-severity threats. You also can create exceptions that allow you to change the response to a specific signature.

After the firewall detects a threat event, you can configure the Vulnerability Protection Profile to perform the actions similar to the anti-spyware response—default, allow, alert, drop, reset-client, reset-server, reset-both, and block IP.

Note that similar to the anti-spyware, when the vulnerability protection action profile is set to reset-both, the associated threat log might display action as reset-server. As discussed earlier, this occurs when the firewall detects the threat at the beginning of a session and presents the client with a 503-block page. Since, the block place disallows the connection, only the server-side connection is reset.

### **URL Filtering Profiles**

A URL Filtering Profile is a collection of URL filtering controls that are applied to individual Security policy rules to enforce your web access policy. The firewall comes with a default profile that is configured to block threat-prone categories, such as malware, phishing, and adult content. You can use the default profile in a Security policy, clone it to be used as a starting point for new URL Filtering Profiles, or add a new URL Filtering Profile that has all the categories set to allow for visibility into the traffic on the network. You can then customize the newly added URL Filtering Profiles and add lists of specific websites that should always be blocked or allowed. This information provides more granular control over URL categories. For example, you might want to block social networking sites but allow some websites that are part of the social networking category.

Configure user-credential detection so that users can submit credentials only to the sites in specified URL categories, which reduces the attack surface by preventing credential submission to sites in untrusted categories. If you block all the URL categories in a URL Filtering profile for user credential submission, you don't need to check credentials.

URL Filtering Profiles enable you to monitor and control how users access the web over HTTP and HTTPS. The firewall comes with a default profile that is configured to block websites, such as known malware sites, phishing sites, and adult content sites. You can use the default profile in a Security policy, clone it to be used as a starting point for new URL Filtering Profiles, or add a new URL Filtering Profile that will have all the categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL Filtering profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories.

With the advanced URL filtering subscription, Inline Categorization enables real-time analysis of URL traffic by using firewall-based or cloud-based ML models, to detect and prevent malicious phishing variants and JavaScript exploits from entering the network.

## Data Filtering Profiles

Data Filtering Profiles prevent sensitive information, such as credit card numbers or Social Security numbers, from leaving a protected network. The Data Filtering Profile also allows you to filter on keywords, such as a sensitive project name or the word “confidential.” It focuses your profile on the desired file types to reduce false-positives. For example, you may want to search only Word documents or Excel spreadsheets and to only scan web-browsing traffic or FTP. There are now a number of predefined patterns to help you easily configure common filtering requirements. These include Social Security numbers, NHI Identification Numbers, credit cards, and more.

You can create custom data pattern objects and attach them to a Data Filtering Profile to define the type of information you want to filter. Create data pattern objects based on the following:

- **Predefined patterns:** Filter for a list of over 20 predefined patterns, including Social Security numbers, national identity numbers from various nations, credit cards, and other useful patterns.
- **Regular expressions:** Filter for a string of characters.
- **File properties:** Filter for file properties and values based on file type.

**Note:** If you are using a third-party, endpoint DLP solution for populating file properties to indicate sensitive content, this option enables the firewall to enforce your DLP policy.

## File Blocking Profiles

The firewall uses File Blocking Profiles to block specified file types over specified applications and in the specified session flow direction (inbound/outbound/both). You can set the profile to alert or block on upload or download, and you can specify which applications will be subject to the File Blocking Profile. You also can configure custom response pages to appear when a user attempts to download the specified file type. This response page allows the user to pause to consider whether to continue and download a file.

You can define your own custom File Blocking Profiles or choose one of the following predefined profiles when applying file blocking to a Security policy rule. The predefined profiles, which are available with content release version 653 and later, allow you to enable best practice file-blocking settings quickly:

- **Basic file blocking:** Attach this profile to the Security policy rules that allow traffic to and from less-sensitive applications to block the files that are commonly included in malware attack campaigns or that have no real use case for upload or download. This profile blocks upload and download of PE files (.scr, .cpl, .dll, .ocx, .pif, .exe), Java files (.class, .jar), Help files (.chm, .hlp), and other potentially malicious file types, including .vbe, .hta, .wsf, .torrent, .7z, .rar, and .bat. The profile also prompts users to acknowledge when they attempt to download encrypted-rar or encrypted-zip files. This rule alerts all the other file types to give you complete visibility into all the file types entering and leaving your network.
- **Strict file blocking:** Use this stricter profile on the Security policy rules, which allow access to the most sensitive applications. This profile blocks the same file types as the basic file blocking profile, and also the flash, .tar, multi-level encoding, .cab, .msi, encrypted-rar, and encrypted-zip files.

Configure a File Blocking Profile with the following actions:

- **alert:** After the specified file type is detected, a log is generated in the Data Filtering log.
- **block:** After the specified file type is detected, the file is blocked and a customizable block page is presented to the user. A log also is generated in the Data Filtering log.
- **continue:** After the specified file type is detected, a customizable response page is presented to the user. The user can click through the page to download the file. A log is also generated in the Data Filtering log. This type of forwarding action requires user interaction and is therefore only applicable for web traffic.

### WildFire Analysis Profiles

Use a WildFire Analysis Profile to enable the firewall to forward unknown files or email links for WildFire analysis. This detection is for zero-day threats contained in files. The firewall's anti-virus threat detection finds known viruses based on local resources. Specify files to be forwarded for analysis based on application, file type, and transmission direction (upload or download). Files matched to the profile rule are forwarded to either the WildFire public cloud or the WildFire private cloud (hosted with a WF-500 appliance), depending on the analysis location defined for the rule. If a profile rule is set to forward files to the WildFire public cloud, the firewall also forwards files that match the existing antivirus signatures in addition to unknown files.

You also can use WildFire Analysis Profiles to set up a WildFire hybrid cloud deployment. If you are using a WildFire appliance to analyze sensitive files (such as PDFs) locally, you can specify less-sensitive file types (such as PE files) or file types that are not supported for WildFire appliance analysis (such as APKs) to be analyzed by the WildFire public cloud. Using both WildFire appliance and WildFire cloud for analysis allows you to benefit from a prompt verdict for the files that have already been processed by the cloud and for files that are not supported for appliance analysis. Doing so also frees the appliance capacity to process sensitive content.

The WildFire cloud can scan your organization's files by using an appropriately configured WildFire Analysis Profile. A profile includes the match conditions describing the file characteristics you want to forward to WildFire for analysis. When any files matching these conditions are transferred through your firewall, a copy is sent to WildFire for analysis.

**Note:** Files are *not* quarantined pending WildFire evaluation. In cases of positive malware findings, the security engineer must use the information collected on the firewall and by WildFire to locate the file internally for remediation.

WildFire Analysis Profiles indicate which files are to be forwarded according to system-wide WildFire configuration settings. WildFire typically renders a verdict on a file within 5 to 10 minutes of receipt.

WildFire analysis results in a detailed report, including all of the aspects of the original file and the contained malware. This report is a valuable tool that describes the exact nature of the detected threat.

## DoS Protection Profiles

DoS Protection profiles provide detailed control for DoS Protection policy rules. DoS Protection profiles 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. The Palo Alto Networks firewalls support the following two DoS-protection mechanisms:

- **Flood protection:** Detects and prevents attacks in which the network is flooded with packets, which results in too many half-open sessions or services being unable to respond to each request. In this case, the source address of the attack is usually spoofed.
- **Resource protection:** Detects and prevents session exhaustion attacks. In this type of attack, many hosts (bots) are used to establish as many fully established sessions as possible for consuming all of a system's resources.

You can enable both types of protection mechanisms in a single DoS Protection profile.

The DoS Protection profile is used to specify the type of action to take and the details on matching criteria for the DoS policy. The DoS Protection profile defines threshold settings for the synchronize (SYN), UDP, and Internet Control Message Protocol (ICMP) floods; enables resource protection; and defines the maximum number of concurrent connections. After configuring the DoS Protection profile, you attach it to a DoS policy rule.

When you configure DoS protection, you should analyze your environment to set the correct thresholds based on your actual traffic rather than the default values provided.

### 2.2.2 Relationship between URL filtering and credential theft prevention

#### Phishing Prevention Overview

The Palo Alto Networks URL filtering solution complements App-ID by controlling access to web (HTTP and HTTPS) traffic and protecting your network from attack.

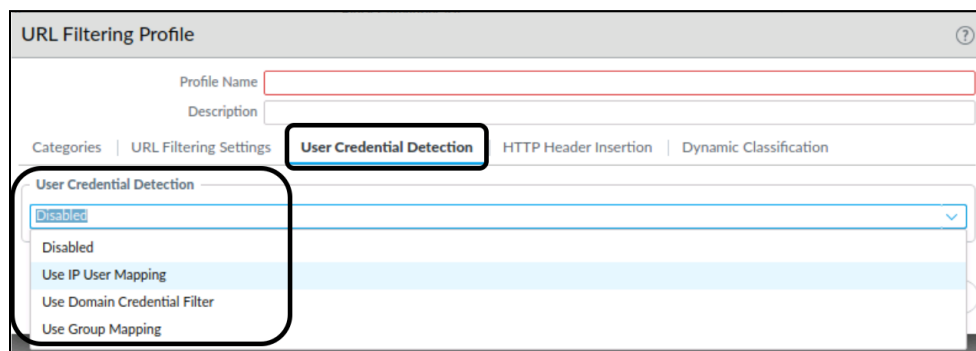
With URL filtering enabled, all of the web traffic is compared against the URL filtering database, which contains a list of millions of categorized websites. You can use these URL categories as match criteria to enforce the Security policy, safely enable web access, and control the traffic that traverses your network. You can also use URL filtering to enforce safe search settings for your users and prevent credential phishing based on URL category.

Credential phishing prevention works by scanning username and password submissions to websites and comparing those submissions against valid corporate credentials. You can choose which websites you want to allow or block corporate credential submissions based on the URL category of the website. When the firewall detects a user credential being transmitted to a site in a category you have restricted, it either displays a block response page that prevents the user from submitting credentials or presents a continue page that warns users against submitting credentials to the sites classified in certain URL categories. The firewall still allows the user to continue with the credential submission. You can customize these block pages to educate users against reusing corporate credentials, even on legitimate, non-phishing sites.

## Credential Detection

Before you configure credential phishing protection, decide which method you want the firewall to use to identify credentials. Each method requires the configuration of User-ID technology. The *IP address-to-username mapping* and *group mapping* methods check for valid username submissions only. In these cases, the firewall blocks or allows the submission based on your settings, regardless of the accompanying password submitted. The *domain credential filter* method checks for valid usernames and passwords submitted to a web page, as follows:

- **IP address-to-username mapping (using PAN-OS-integrated agent):** The firewall uses IP-address-to-user mappings that User-ID collects to check if a username submitted to a webpage matches the username of the logged-in user.
- **Group mapping (using PAN-OS integrated agent):** The User-ID agent collects group mapping information from a directory server and retrieves a list of groups and corresponding group members. It compares the usernames submitted to a webpage against the group member usernames.
- **Domain credential filter (using Windows-based agent):** The User-ID agent is installed on a Read-Only Domain Controller. The User-ID agent collects password hashes that correspond to users for whom you want to enable credential detection, and it sends these mappings to the firewall. The firewall then checks if the source IP address of a session matches a username and if the password submitted to the web page belongs to that username. With this mode, the firewall only blocks or alerts on the submission when the submitted password matches a user password.



## Category Selection for Enforcement

After the detection method is chosen for the URL Filtering Profile, the enforcement action must be chosen for each appropriate browsing category. Custom categories can be created when flexibility is required for identifying specific category members. For each category, select how you want to treat user credential submissions from the following:

- **alert:** Allow users to submit credentials to the website, but generate a URL Filtering log each time a user submits credentials to the sites in this URL category.
- **allow:** (default) Allow users to submit credentials to the website.
- **block:** Block users from submitting credentials to the website. When a user tries to submit credentials, the firewall displays the Anti Phishing Block page, which prevents the credential submission.
- **continue:** The firewall displays the Anti Phishing Continue page response page when a user attempts to submit credentials. Users must select **Continue** on the response page to continue with the submission.

When the firewall detects a user attempting to submit credentials to a site in a category that you have restricted, it either displays a block response page that prevents the user from submitting credentials or presents a continue page that warns the user against submitting credentials to sites classified in certain URL categories. The firewall still allows the user to continue with the credential submission. You can customize these block pages to educate users against reusing corporate credentials, even on legitimate, non-phishing sites.

### 2.2.3 Use of username and domain name in HTTP header insertion

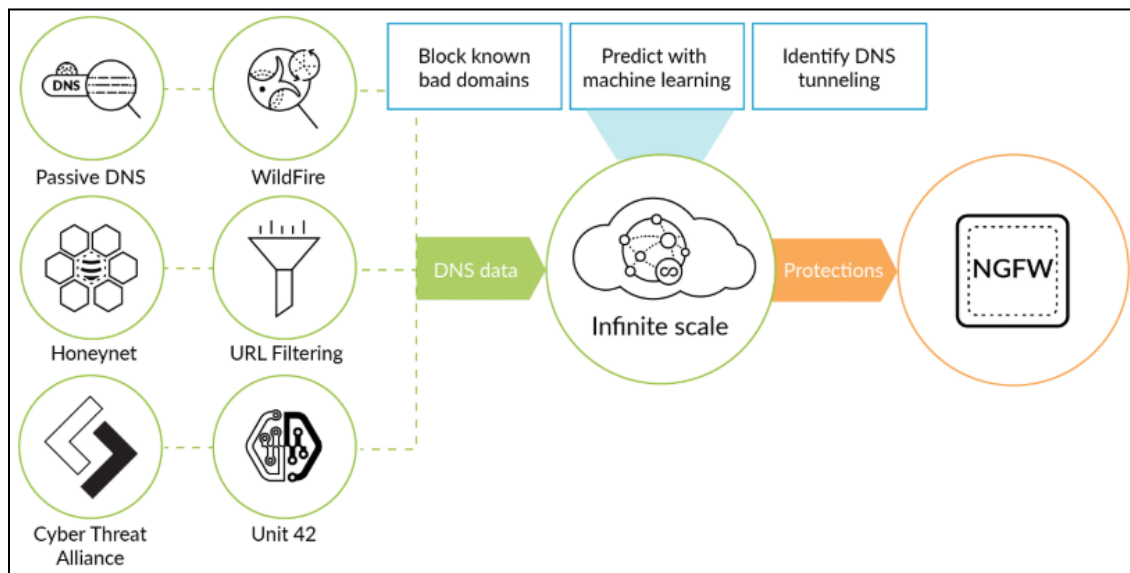
The firewall supports header insertion for HTTP/1.x traffic only. The firewall does not support header insertion for HTTP/2 traffic. You can create insertion entries based on a predefined HTTP header insertion type, or you can create your own custom type. Header insertion is performed for custom HTTP headers. You can also insert standard HTTP headers.

HTTP header insertion can only be performed by using the following methods:

- GET
- POST
- PUT
- HEAD

### 2.2.4 DNS Security

DNS Security allows you to apply predictive analytics, ML, and automation to block the attacks that use DNS. Tight integration with NGFW gives you automated protection and eliminates the need for independent tools. Now you can rapidly predict and prevent malicious domains, neutralize threats hidden in DNS tunneling, and apply automation to quickly find and contain infected devices. The following illustration depicts DNS Security sources, intermediate processing of source data, and ultimate delivery to a firewall:



## 2.2.5 How to tune or add exceptions to a Security Profile

Palo Alto Networks defines a recommended default action (such as block or alert) for threat signatures. You can use a threat ID to exclude a threat signature from enforcement or modify the action the firewall enforces for that threat signature. For example, you can modify the action for the threat signatures that are triggering false-positives on your network.

Configure threat exceptions for antivirus, vulnerability, spyware, and DNS signatures to change firewall enforcement for a threat. However, before you begin, make sure the firewall is detecting and enforcing threats based on the default signature settings by ensuring the following:

- Get the latest Antivirus, Threats and Applications, and WildFire signature updates.
- Set up Antivirus, Anti-Spyware, and Vulnerability Protection subscriptions, and apply these Security Profiles to your Security policy.

**Step 1:** Exclude antivirus signatures from enforcement.

- Select **Objects > Security Profiles > Antivirus**.
- Add or modify an existing Antivirus Profile from which you want to exclude a threat signature, and select **Signature Exceptions**.
- Add the **Threat ID** for the threat signature you want to exclude from enforcement.

The screenshot shows the 'Signature Exceptions' configuration page for 'WildFire Inline ML'. At the top, there are tabs for 'Action', 'Signature Exceptions', and 'WildFire Inline ML'. Below the tabs is a search bar with a magnifying glass icon and a '1 item' indicator. The main content area contains a table with the following data:

THREAT ID ^	THREAT NAME	
280647	JS/Exploit.pdfka.os	<input type="checkbox"/>

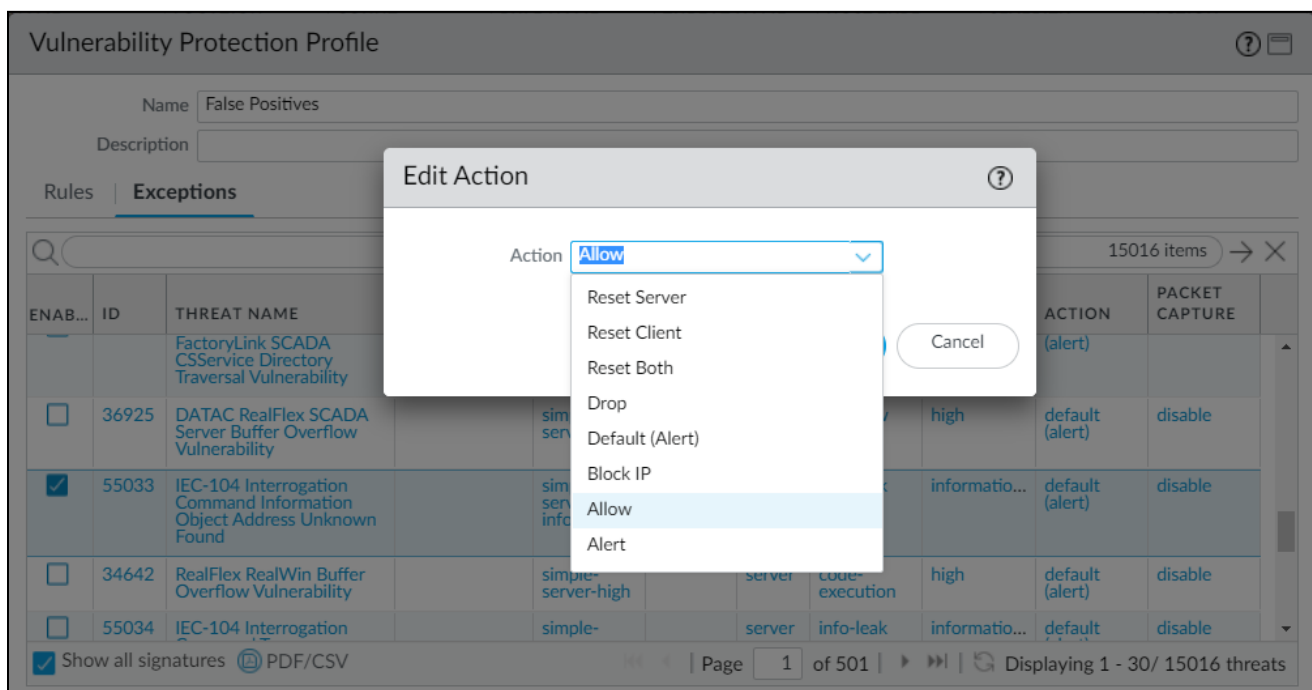
At the bottom of the table, there is a search bar with 'Threat ID 280647' entered, an 'Add' button with a plus sign, and a 'PDF/CSV' button.

- Click **OK** to save the Antivirus Profile.



**Step 2:** Modify enforcement for vulnerability and spyware signatures. (This does not include DNS signatures; skip to the next option to modify enforcement for DNS signatures, which are a type of spyware signature.)

- Select **Objects > Security Profiles > Anti-Spyware** or **Objects > Security Profiles > Vulnerability Protection**.
- Add or modify an existing Anti-Spyware Profile or Vulnerability Protection Profile from which you want to exclude the threat signature. Then, select either **Signature Exceptions for Anti-Spyware Protection profiles** or **Exceptions for Vulnerability Protection profiles**.
- Show all the signatures, and then filter to select the signature for which you want to modify enforcement rules.
- Check the box under the **Enable** column for the signature for which you want to modify enforcement.
- Select the **Action** you want the firewall to enforce for this threat signature.

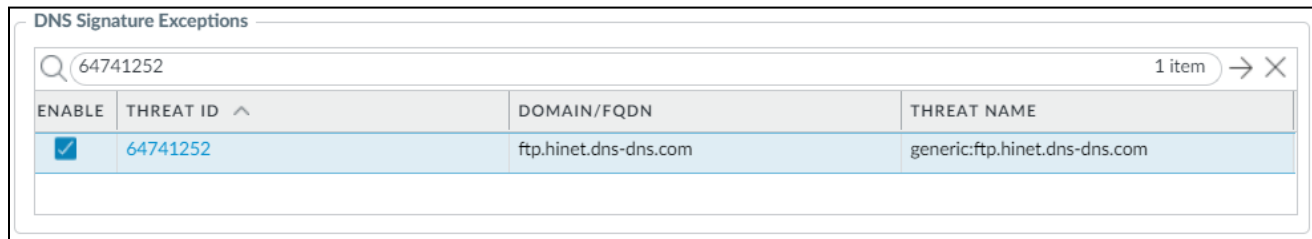


- For signatures that you want to exclude from enforcement because they trigger false-positives, set the **Action** to **Allow**.
- Click **OK** to save the new or modified Anti-Spyware Profile or Vulnerability Protection Profile.

### Step 3: Modify enforcement for DNS signatures.

By default, the DNS lookups for malicious hostnames that the DNS signatures detect are sinkholed.

- Select **Objects > Security Profiles > Anti-Spyware**.
- Add or modify the Anti-Spyware Profile from which you want to exclude the threat signature, and then select **DNS Exceptions**.
- Search for the **DNS Threat ID** for the DNS signature that you want to exclude from enforcement, and then select the box of the applicable signature.



ENABLE	THREAT ID ^	DOMAIN/FQDN	THREAT NAME
<input checked="" type="checkbox"/>	64741252	ftp.hinet.dns-dns.com	generic:ftp.hinet.dns-dns.com

- Click **OK** to save the new or modified Anti-Spyware Profile.

## 2.2.6 Compare and contrast threat prevention and advanced threat prevention

### Threat Prevention

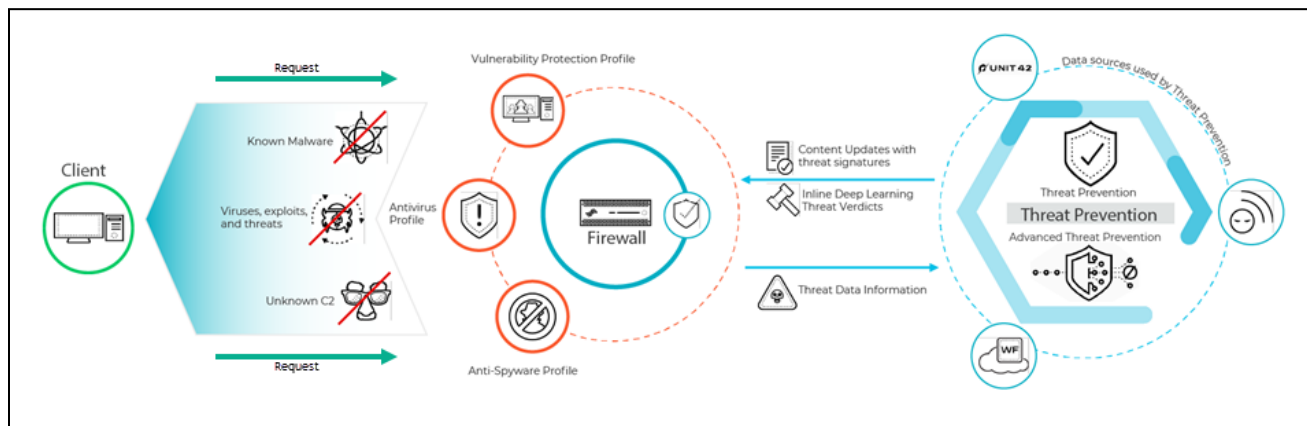
Threat Prevention is an IPS solution that can detect and block malware, vulnerability exploits, and C2 across all of the ports and protocols. It uses a multilayered prevention system with components operating on the firewall and in the cloud. The Threat Prevention cloud operates a multitude of detection services by using the combined threat data from the Palo Alto Networks services to create signatures, each possessing specific identifiable patterns that are used by the firewall to enforce security policies when matching threats and malicious behaviors are detected. These signatures are categorized based on the threat type and are assigned unique identifier numbers. To detect threats that correspond with these signatures, the firewall operates analysis engines that inspect and classify the network traffic exhibiting anomalous traits.

### Advanced Threat Prevention

Advanced Threat Prevention is a cloud-delivered security service that works in conjunction with the existing Threat Prevention license to deliver protections for advanced and evasive C2 threats. Advanced Threat Prevention allows you to prevent unknown threats by using real-time traffic inspection and inline detectors. These deep learning, ML-based detection engines in the Advanced Threat Prevention cloud analyze traffic for advanced C2 and spyware threats to protect users against zero-day threats. By operating cloud-based detection engines, you can access a wide array of detection mechanisms that are updated and deployed automatically without requiring the user to download update packages or operate process-intensive, firewall-based analyzers, which can sap resources. The cloud-based detection engine logic is continuously monitored and updated using C2 traffic datasets from WildFire, with additional support from the Palo Alto Networks threat researchers who provide human intervention for highly accurate detection enhancements. Advanced Threat Prevention deep learning engines support the analysis of C2-based threats over HTTP, HTTP2, SSL, unknown-UDP, and unknown-TCP applications. Additional analysis models are delivered through content updates; however, enhancements to existing models are performed as a

cloud-side update, requiring no firewall update. Advanced Threat Prevention is enabled and configured under inline cloud analysis in the Anti-Spyware Profile.

In addition to the signature-based detection mechanism, Advanced Threat Prevention provides a complementary inline detection system to prevent unknown and evasive C2 threats. The Advanced Threat Prevention cloud operates deep learning models that enable inline analysis on the firewall on a per-request basis to prevent zero-day threats from entering the network.



The threat signatures used by the firewall are broadly categorized into three types: antivirus, anti-spyware, and vulnerability. These types are used by Security Profiles to enforce the user-defined policy, as follows:

- Antivirus signatures detect various types of malware and viruses, including worms, Trojan horses, and spyware downloads.
- Anti-spyware signatures detect C2 spyware on compromised hosts that try to phone-home or beacon out to an external C2 server.
- Vulnerability signatures detect exploit system vulnerabilities.

Signatures have a default severity level with an associated default action; for example, in the case of a highly malicious threat, a setting of *reset both*. This setting is based on security recommendations from Palo Alto Networks. In deployments that use specialized internal applications or third-party intelligence feeds with open source SNORT and Suricata rules, custom signatures can be created for purpose-built protection. Firewalls receive signature updates from two update packages: the daily antivirus content update and the weekly application and threats content update. The antivirus content updates include the antivirus signatures and DNS (C2) signatures used by the Antivirus and Anti-Spyware Profiles, respectively. The Applications and Threats content updates include the vulnerability and anti-spyware signatures used by the Vulnerability and Anti-Spyware Profiles, respectively. The update packages also include content that is leveraged by other services and subfunctions.

## 2.2.7 Compare and contrast URL Filtering and Advanced URL Filtering

### URL Filtering

- URL Filtering allows you to protect your organization against web-based threats, such as phishing, malware, and C2. Inline ML instantly identifies and prevents new and unknown malicious websites before they can be accessed by users. Web Security rules are an extension of the organization's NGFW policy, thus reducing complexity by giving you a single policy set to manage.

URL Filtering provides the following benefits:

- Reduces infection risk from dangerous websites and protects users and data from malware and credential-phishing pages
- Protects across the attack lifecycle through integration with WildFire and the cybersecurity portfolio
- Retains protections synchronized with the latest threat intelligence through the Palo Alto Networks cloud-based URL categorization for phishing, malware, and undesired content
- Provides full visibility and threat inspection into normally opaque web traffic through granular control over SSL decryption

### Advanced URL Filtering

Built in the cloud, Advanced URL Filtering is a subscription service that works natively with the Palo Alto Networks NGFW to secure your network against web-based threats, such as phishing, malware, and C2. Advanced URL Filtering uses ML to analyze URLs in real time and classify them into benign or malicious categories, which you can easily build into your NGFW policy for total control of web traffic. These categories trigger complementary capabilities across the NGFW platform, enabling additional layers of protection, such as targeted SSL decryption and advanced logging. Alongside its own analysis, Advanced URL Filtering uses shared threat information from WildFire, Palo Alto Networks' industry-leading malware prevention service, and other sources to automatically update protections against malicious sites. Advanced URL Filtering delivers:

- Superior protection against web-based attacks with the combined power of our URL database and cloud-delivered web security engine powered by ML that categorizes and blocks new malicious URLs in real time, even when content is cloaked from crawlers. Advanced URL Filtering prevents 40 percent more threats than traditional web-filtering databases.
- Industry-leading phishing protections that tackle the most common causes of breaches.
- Total control of web traffic through fine-grained controls and policy settings that enable you to automate security actions based on users, risk ratings, and content categories.
- Maximum operational efficiency by enabling web protection through the Palo Alto Networks platform.

## 2.2.8 References

Security Profiles,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/policy/security-profiles>

Customize the URL Filtering Response Pages,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/url-filtering/customize-the-url-filtering-response-pages>

URL Filtering,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/url-filtering>

Configure URL Filtering,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/url-filtering/configure-url-filtering>

WildFire Analysis Reports—Close Up,

<https://docs.paloaltonetworks.com/wildfire/11-0/wildfire-admin/monitor-wildfire-activity/wildfire-analysis-reports-close-up>

Objects > Security Profiles > WildFire Analysis,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/objects/objects-security-profiles-wildfire-analysis>

WildFire Administrator's Guide,

<https://docs.paloaltonetworks.com/wildfire/11-0/wildfire-admin>

WildFire Subscription,

<https://docs.paloaltonetworks.com/wildfire/11-0/wildfire-admin/wildfire-overview/wildfire-subscription>

Take a Threat Packet Capture,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/take-packet-captures/take-a-threat-packet-capture>

Enable Data Capture for Data Filtering and Manage Data Protection Password,

<https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000ClvCAC>

Threat Prevention,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/threat-prevention>

Create Threat Exceptions,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/threat-prevention/create-threat-exceptions>

## 2.3 Configure zone protections, packet buffer protection, and DoS protection

### 2.3.1 Customized values versus default settings

#### Flood Protection

A Zone Protection profile with flood protection configured defends an entire ingress zone against SYN, ICMP, ICMPv6, UDP, and other IP flood attacks. The firewall measures the aggregate amount of each flood type entering the zone in new connections per second (CPS) and compares the totals to the thresholds you configure in the Zone Protection profile.

For each flood type, you set three thresholds for new CPS entering the zone, and you can set a drop Action for SYN floods. If you know the baseline CPS rates for the zone, use these guidelines to set the initial thresholds, and then monitor and adjust the thresholds as necessary.

- **Alarm Rate**—The new CPS threshold to trigger an alarm. Target setting the Alarm Rate to 15-20% above the average CPS rate for the zone so that normal fluctuations don't cause alerts.
- **Activate**—The new CPS threshold to activate the flood protection mechanism and begin dropping new connections. For ICMP, ICMPv6, UDP, and other IP floods, the protection mechanism is Random Early Drop (RED, also known as Random Early Detection). For SYN floods only, you can set the drop Action to SYN Cookies or RED. Target setting the Activate rate to just above the peak CPS rate for the zone to begin mitigating potential floods.
- **Maximum**—The number of connections-per-second to drop incoming packets when RED is the protection mechanism. Target setting the Maximum rate to approximately 80-90% of firewall capacity, taking into account other features that consume firewall resources.

If you don't know the baseline CPS rates for the zone, start by setting the Maximum CPS rate to approximately 80-90% of firewall capacity and use it to derive reasonable flood mitigation alarm and activation rates. Set the Alarm Rate and Activate rate based on the Maximum rate. The default threshold values are high so that activating a Zone Protection profile doesn't unexpectedly drop legitimate traffic. Adjust the thresholds to values appropriate for your network's traffic. The best method for understanding how to set reasonable flood thresholds is to take baseline measurements of average and peak CPS for each flood type to determine the normal traffic conditions for each zone and to understand the capacity of the firewall, including the impact of other resource-consuming features such as decryption. Monitor and adjust the flood thresholds as needed and as your network evolves.

### Reconnaissance Protection

Similar to the military definition of reconnaissance, the network security definition of reconnaissance is when attackers attempt to gain information about your network's vulnerabilities by secretly probing the network to find weaknesses. Reconnaissance activities are often preludes to a network attack. Enable Reconnaissance Protection on all zones to defend against port scans and host sweeps:

- **Port scans** discover open ports on a network. A port scanning tool sends client requests to a range of port numbers on a host, with the goal of locating an active port to exploit in an attack. Zone Protection profiles defend against TCP and UDP port scans.
- **Host sweeps** examine multiple hosts to determine if a specific port is open and vulnerable.

You can use reconnaissance tools for legitimate purposes such as penetration (pen) testing of network security or the strength of a firewall. You can specify up to 20 IP addresses or netmask address objects to exclude from Reconnaissance Protection so that your internal IT department can conduct pen tests to find and fix network vulnerabilities.

## Packet-Based Attack Protection

Packet-based attacks take many forms. Zone Protection profiles check IP, TCP, ICMP, IPv6, and ICMPv6 packet headers and protect a zone by:

- Dropping packets with undesirable characteristics.
- Stripping undesirable options from packets before admitting them to the zone.

Select the drop characteristics for each packet type when you Configure Packet Based Attack Protection. The best practices for each IP protocol are:

- **IP Drop**—Drop Unknown and Malformed packets. Also drop Strict Source Routing and Loose Source Routing because allowing these options allows adversaries to bypass Security policy rules that use the Destination IP address as the matching criteria. For internal zones only, check Spoofed IP Address so only traffic with a source address that matches the firewall routing table can access the zone.
- **TCP Drop**—Retain the default TCP SYN with Data and TCP SYN-ACK with Data drops, drop Mismatched overlapping TCP segment and Split Handshake packets, and strip the TCP Timestamp from packets.
- **ICMP Drop**—There are no standard best practice settings because dropping ICMP packets depends on how you use ICMP (or if you use ICMP). For example, if you want to block ping activity, you can block ICMP Ping ID 0.
- **IPv6 Drop**—If compliance matters, ensure that the firewall drops packets with non-compliant routing headers, extensions, etc.
- **ICMPv6 Drop**—If compliance matters, ensure that the firewall drops certain packets if the packets don't match a Security policy rule.

## Protocol Protection

In a Zone Protection profile, Protocol Protection defends against non-IP protocol-based attacks. Enable Protocol Protection to block or allow non-IP protocols between security zones on a Layer 2 VLAN or on a virtual wire, or between interfaces within a single zone on a Layer 2 VLAN (Layer 3 interfaces and zones drop non-IP protocols so non-IP Protocol Protection doesn't apply). Configure Protocol Protection to reduce security risks and facilitate regulatory compliance by preventing less secure protocols from entering a zone, or an interface in a zone.

Create an **Exclude List** or an **Include List** to configure Protocol Protection for a zone. The **Exclude List** is a block list—the firewall blocks all of the protocols you place in the **Exclude List** and allows all other protocols. The **Include List** is an allow list—the firewall allows only the protocols you specify in the list and blocks all other protocols.

## Ethernet SGT Protection

In a Cisco TrustSec network, a Cisco Identity Services Engine (ISE) assigns a Layer 2 Security Group Tag (SGT) of 16 bits to a user's or endpoint's session. You can create a Zone Protection profile with Ethernet SGT protection when your firewall is part of a Cisco TrustSec network. The firewall can inspect headers with 802.1Q (Ethertype 0x8909) for specific Layer 2 security group tag (SGT) values and drop the packet if the SGT matches the list you configure for the Zone Protection profile attached to the interface. Determine which SGT values you want to deny access to a zone.



### 2.3.2 Classified versus aggregate profile values

You can configure aggregate and classified DoS Protection Profiles and apply one profile or one of each type of profile to DoS Protection Policy Rules when you configure DoS Protection.

- **Aggregate** — Sets thresholds that apply to the entire group of devices specified in a DoS Protection policy rule instead of to each individual device, so one device could receive the majority of the allowed connection traffic. For example, a Max Rate of 20,000 CPS means the total CPS for the group is 20,000 and an individual device can receive up to 20,000 CPS if the other devices don't have connections. Aggregate DoS Protection policies provide another layer of broad protection (after the dedicated DDoS device at the Internet perimeter and Zone Protection profiles) for a particular group of critical devices when you want to apply extra constraints on specific subnets, users, or services.
- **Classified** — Sets flood thresholds that apply to each individual device specified in a DoS Protection policy rule. For example, if you set a Max Rate of 5,000 CPS, each device specified in the rule can accept up to 5,000 CPS before it drops new connections. If you apply a classified DoS Protection policy rule to more than one device, the devices governed by the rule should be similar in terms of capacity and how you want to control their CPS rates because classified thresholds apply to each individual device. Classified profiles protect individual critical resources.

How you configure the **Address (source-ip-only, destination-ip-only, or src-dest-ip-both)** for classified profiles depends on your DoS protection goals, what you are protecting, and whether the protected device(s) are in internet-facing zones.

If you apply both an aggregate and a classified DoS Protection profile to the same DoS Protection policy rule, the firewall applies the aggregate profile first and then applies the classified profile if needed. For example, we protect a group of five web servers with both types of profiles in a DoS Protection policy rule. The aggregate profile configuration drops new connections when the combined total for the group reaches a **Max Rate** of 25,000 CPS. The classified profile configuration drops new connections to any individual web server in the group when it reaches a **Max Rate** of 6,000 CPS. There are three scenarios in which new connection traffic crosses **Max Rate** thresholds:

- The new CPS rate exceeds the aggregate **Max Rate** but doesn't exceed the classified **Max Rate**. In this scenario, the firewall applies the aggregate profile and blocks all new connections for the configured Block Duration.
- The new CPS rate doesn't exceed the aggregate **Max Rate**, but the CPS to one of the web servers exceeds the classified **Max Rate**. In this scenario, the firewall checks the aggregate profile and finds that the rate for the group is less than 25,000 CPS, so the firewall doesn't block new connections based on that. Next, the firewall checks the classified profile and finds that the rate for a particular server exceeds 6,000 CPS. The firewall applies the classified profile and blocks all new connections to that particular server for the configured Block Duration. The other servers in the group are within the classified profile's **Max Rate**, and so their traffic is not affected.



- The new CPS rate exceeds the aggregate **Max Rate** and also exceeds the classified **Max Rate** for one of the web servers. In this scenario, the firewall checks the aggregate profile and finds that the rate for the group exceeds 25,000 CPS, so the firewall blocks new connections to limit the group's total CPS. The firewall then checks the classified profile and finds that the rate for a particular server exceeds 6,000 CPS (so the aggregate profile enforced the group's combined limit but wasn't enough to protect this particular server). The firewall applies the classified profile and blocks all new connections to that particular server for the configured Block Duration. The other servers in the group are within the classified profile's **Max Rate**, so their traffic is not affected.

### 2.3.3 Layer 3 and Layer 4 header inspection

When L3 & L4 header inspection is enabled globally, the firewall can detect and prevent vulnerabilities within the supported protocols (IP/IPv6, ICMP/ICMPv6, TCP and UDP) and log and/or block packets that match the user-specified custom rules. Additionally, you must Enable Net Inspection (**NetworkZones**) for each security zone by using header inspection custom rules.

You can add, delete, and clone existing rules, as well as define the precedence and operational status of the custom rules as evaluated by the Zone Protection profile.

After you configure L3 & L4 header inspection in a Zone Protection profile, apply the profile to an ingress security zone.

### 2.3.4 References

- Zone Protection Profiles,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/zone-protection-and-dos-protection/zone-defense/zone-protection-profiles>
- Classified versus Aggregate DoS Protection,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/zone-protection-and-dos-protection/zone-defense/dos-protection-profiles-and-policy-rules/classified-versus-aggregate-dos-protection>
- L3 & L4 Header Inspection,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/network/network-network-profiles/network-network-profiles-zone-protection/l3-l4-header-inspection>

## 2.4 Design the deployment configuration of a Palo Alto Networks firewall

### 2.4.1 Advanced high availability (HA) deployments

You can set up the firewalls in an HA pair in one of two modes:

- **Active/Passive** — One firewall actively manages traffic while the other is synchronized and ready to transition to the active state, should a failure occur. In this mode, both firewalls share the same configuration settings and one actively manages traffic until a path, link, system, or network failure occurs. When the active firewall fails, the passive firewall transitions to the active state and takes over seamlessly and enforces the same policies to maintain network security. Active/passive HA is supported in the virtual wire, Layer 2, and Layer 3 deployments.
- **Active/Active** — Both firewalls in the pair are active, process traffic, and work synchronously to handle session setup and session ownership. Both firewalls individually maintain session tables and routing tables and synchronize with each other. Active/active HA is supported in virtual wire and Layer 3 deployments.  
In active/active HA mode, the firewall does not support the DHCP client. Furthermore, only the active-primary firewall can function as a DHCP Relay. If the active-secondary firewall receives DHCP broadcast packets, it drops them.

When deciding whether to use active/passive or active/active mode, consider the following differences:

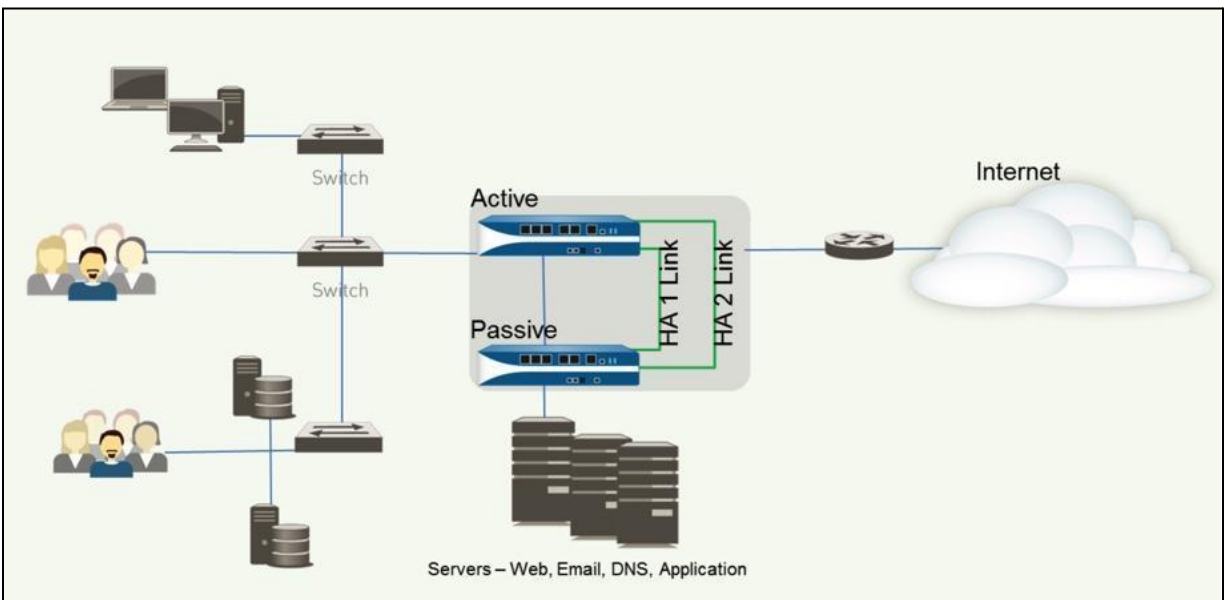
- Active/passive mode has simplicity of design; it is significantly easier to troubleshoot the routing and traffic flow issues in active/passive mode. Active/passive mode supports a Layer 2 deployment while active/active mode does not.
- Active/active mode requires advanced design concepts that can result in more complex networks. Depending on how you implement active/active HA, it might require additional configuration, such as activating networking protocols on both firewalls, replicating NAT pools, and deploying floating IP addresses to provide proper failover. The firewalls use additional concepts of session owner and session setup to perform Layer 7 content inspection because both firewalls actively process traffic. Active/active mode is recommended if each firewall needs its own routing instances and you require full, real-time redundancy out of both firewalls all the time. Active/active mode has faster failover and can handle peak traffic flows better than active/passive mode because both firewalls actively process traffic.

## 2.4.2 HA Pair

### Active/Passive HA Pair

To set up an active/passive HA pair as shown below, complete the following high-level tasks:

- Connect the HA ports.
- Enable ping on the management port.
- Set the HA mode and group ID.
- Set up the control link connection.
- (Optional) Enable encryption for the control link connection.
- Set up the backup control link connection.
- Set up the data link connection (HA2) and the backup HA2 connection.
- Enable heartbeat backup.
- Set the device priority and enable preemption.
- (Optional) Modify the HA timers.
- (Optional) Modify the link status of the HA ports on the passive device.
- Enable HA.
- (Optional) Enable LACP and LLDP Pre-Negotiation for active/passive HA.
- Verify that the firewalls are paired.



## Active/Active HA Pair

The basic workflow for configuring firewalls in an active/active pair is as follows:

- Determine the active/active use case.
- Connect the HA ports.
- Enable ping on the management port.
- Enable active/active HA and set the group ID.
- Set the Device ID, enable synchronization, and identify the control link on the peer firewall.
- Enable heartbeat backup.
- (Optional) Modify the HA timers.
- Set up the control link connection.
- (Optional) Enable encryption for the control link connection.
- Set up the backup control link connection.
- Set up the data link connection (HA2) and the backup HA2 connection.
- Configure the HA3 link for packet forwarding.
- (Optional) Modify the Tentative Hold time.
- Configure Session Owner and Session Setup.
- Configure an HA virtual address.
- Configure the floating IP address.
- Configure ARP load-sharing.
- Define HA failover conditions.
- Commit the configuration.

### 2.4.3 Zero-Touch Provisioning

To set up your firewall for Zero-Touch Provisioning (ZTP), perform the following using Panorama:

- Select **Panorama > Plugins to Download**. Install the most recent version of the ZTP plugin.
- Install the Panorama device certificate.
- Register Panorama with the ZTP service.
- Create a default device group and template to connect the ZTP firewalls to Panorama.
- Select **Panorama > Zero Touch Provisioning** and sync Panorama with the ZTP service.
- Set up the ZTP installer administrative account.
- Add ZTP firewalls to Panorama.

### 2.4.4 Bootstrapping

#### Bootstrapping

All Palo Alto Networks firewalls can automatically configure themselves during first boot by using the bootstrapping feature. This feature provisions a specifically prepared storage volume (USB for physical appliances or storage accounts for VM-Series firewalls) containing configuration data, licenses, dynamic updates, and PAN-OS updates. These are all applied automatically during the boot process.

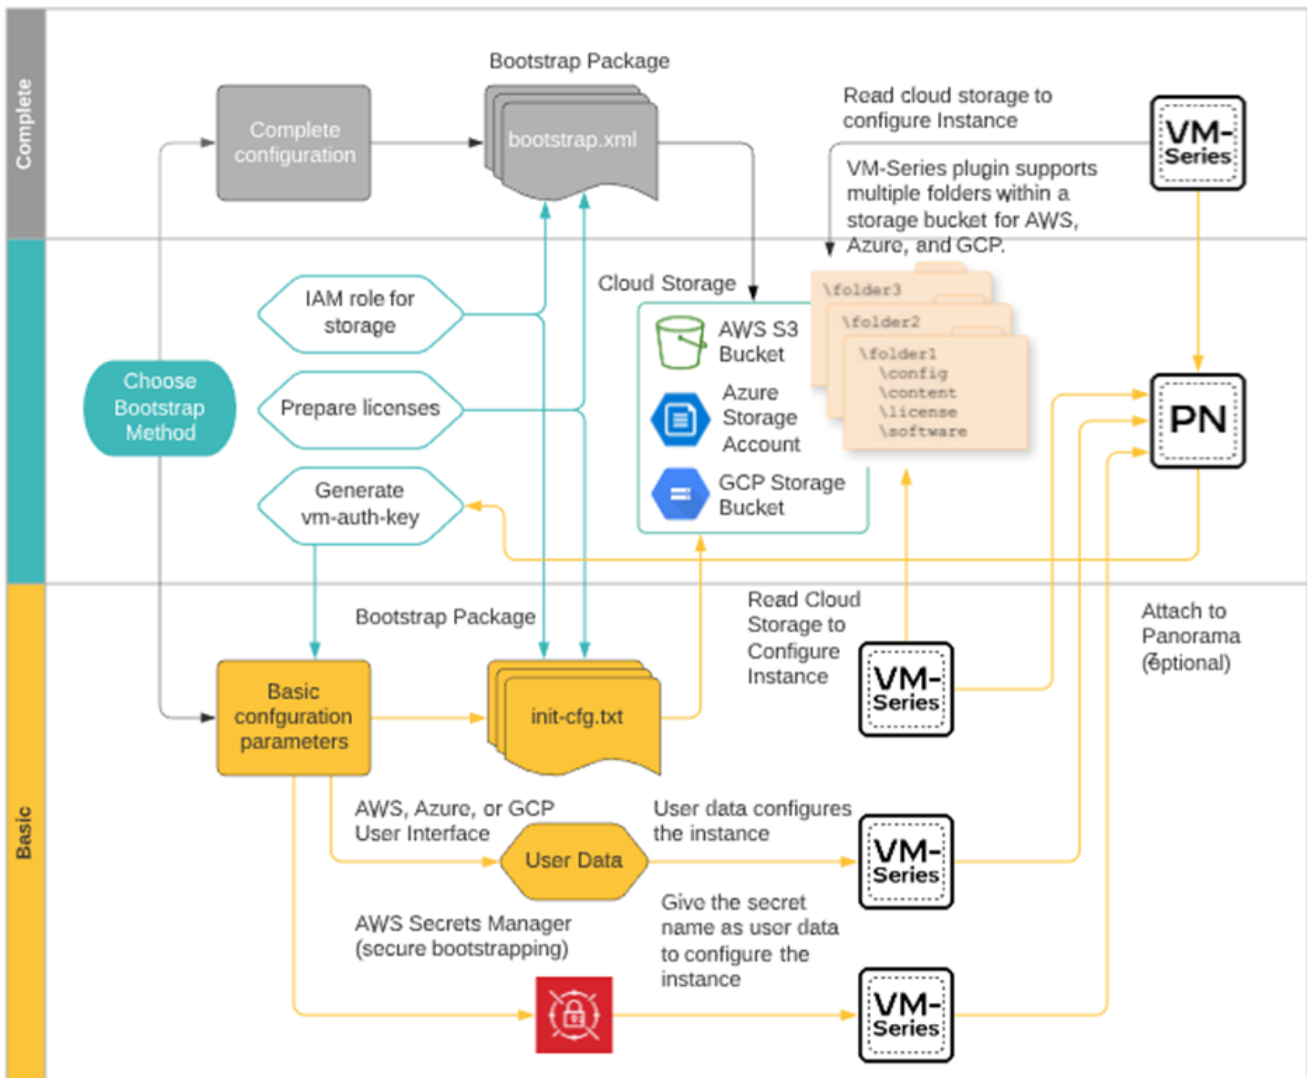
#### VM-Series Bootstrapping

Bootstrapping enables you to create a repeatable and streamlined process of deploying new VM-Series firewalls on the network. It allows you to create a package with the model configuration for the network and then use that package to deploy the VM-Series firewalls anywhere. You can

bootstrap the VM-Series firewall off an external device (such as a virtual disk, a virtual CD-ROM, or an Amazon Web Services S3 or Google Cloud bucket) to configure and license the VM-Series firewall. You can bootstrap the firewall with a basic initial configuration and licenses so that the firewall can register with Panorama and then retrieve its full configuration from Panorama. You can also bootstrap the complete configuration so that the firewall is fully configured on bootup.

The workflow to set up bootstrapping on a VM-Series firewall as shown in the image below is as follows:

- Reset the firewall to factory default settings.
- Choose a bootstrap method.
- (Optional) Generate the VM auth key on Panorama.
- Prepare the licenses for bootstrapping.
- Prepare the bootstrap package and save the bootstrap package in the appropriate delivery format for your hypervisor.
- Bootstrap the VM-Series firewall.
- Verify bootstrap completion.



## Bootstrap Package

The bootstrap process is initiated only when the firewall starts up in a factory default state. After you attach the virtual disk, virtual CD-ROM, or storage bucket to the firewall, the firewall scans for a bootstrap package. If one exists, the firewall uses the settings defined in the bootstrap package. If you have included a Panorama server IP address in the file, the firewall connects with Panorama. If the firewall has Internet connectivity, it contacts the licensing server to update the universally unique identifier (UUID) and obtain the license keys and subscriptions. The firewall is then added as an asset in the Palo Alto Networks Support Portal. If the firewall does not have internet connectivity, it either uses the license keys that you included in the bootstrap package or connects to Panorama, which retrieves the appropriate licenses and deploys them to the managed firewalls.

The bootstrap package that you create must include the `/config`, `/license`, `/software`, and `/content` folders, even if empty, as follows:

- **/config folder:** This folder contains the configuration files. The folder can hold two files, `init-cfg.txt` and `bootstrap.xml`.

**Note:** If you intend to pre-register the VM-Series firewalls with Panorama with bootstrapping, you must generate a VM authorization key on Panorama and include the generated key in the `init-cfg` file.

- **/license folder:** This folder contains the license keys or authorization codes for the licenses and subscriptions that you intend to activate on the firewalls. If the firewall does not have internet connectivity, you must either manually obtain the license keys from the Palo Alto Networks Support Portal or use the Licensing API to obtain the keys and then save each key in this folder.

**Note:** You must include an authorization code bundle instead of individual authorization codes so that the firewall or orchestration service can simultaneously fetch all the license keys associated with a firewall. If you use individual authorization codes instead of a bundle, the firewall will retrieve only the license key for the first authorization code included in the file.

- **/software folder:** This folder contains the software images that are required to upgrade a newly provisioned VM-Series firewall to the desired PAN-OS version for the network. You must include all intermediate software versions between the Open Virtualization Format version and the final PAN-OS software version to which you want to upgrade the VM-Series firewall.
- **/content folder:** This folder contains the Applications and Threats updates and WildFire updates for the valid subscriptions on the VM-Series firewall. You must include the minimum content versions that are required for the desired PAN-OS version. Without the minimum required content version associated with the PAN-OS version, the VM-Series firewall cannot complete the software upgrade.
- **/plugins folder:** This optional folder contains a single VM-Series plugin image.

## 2.4.5 References

HA Modes,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/ha-concepts/ha-modes>

Configure Active/Passive HA,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/set-up-activepassive-ha/configure-activepassive-ha>

Configure Active/Active HA,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/set-up-activeactive-ha/configure-activeactive-ha>

Product Summary Specs sheet,

<https://www.paloaltonetworks.com/resources/datasheets/product-summary-specsheet>

Getting Started,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/getting-started>

Internet Gateway Best Practice Security Policy,

<https://docs.paloaltonetworks.com/best-practices/11-0/internet-gateway-best-practices>

Best Practices,

<https://docs.paloaltonetworks.com/best-practices>

Dynamic Content Updates,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-upgrade/software-and-content-updates/dynamic-content-updates>

Bootstrap the VM-Series Firewall,

<https://docs.paloaltonetworks.com/vm-series/11-0/vm-series-deployment/bootstrap-the-vm-series-firewall>

## 2.5 Configure authorization, authentication, and device access

### 2.5.1 Role-based access control for authorization

#### Administrative Accounts and Roles

Administrators can configure, manage, and monitor Palo Alto Networks firewalls and Panorama by using the web interface, the command line interface (CLI), and the XML API management interface. You can customize role-based administrative access to the management interfaces to delegate specific tasks or permissions to certain administrators.

Administrative accounts specify roles and authentication methods for the administrators of Palo Alto Networks firewalls and Panorama. Each device has a predefined default administrative account (admin) that provides full read-write access (also known as superuser access) to the firewall. Other administrative accounts can be created as needed.

Administrator accounts are configured based on the security requirements of your organization, any existing authentication services that your network uses, and the required administrative roles. A role defines the type of system access that is available to an administrator. You can define and restrict access as broadly or as granularly as required, depending on the security requirements of your organization. For example, you might decide that a data center administrator can have access to all of the device and networking configurations but a security administrator can control only the

Security policy definitions, while other key individuals can have limited CLI or XML API access. The role types are as follows:

**Dynamic roles:** These are built-in roles that provide access to Panorama and managed firewalls. After new features are added, the firewall and Panorama automatically update the definitions of dynamic roles; you do not need to update them manually. The following table lists the access privileges associated with dynamic roles.

DYNAMIC ROLE	PRIVILEGES
Superuser	Full read-write access to Panorama
Superuser (read-only)	Read-only access to Panorama
Panorama administrator	Full access to Panorama except for the following actions: <ul style="list-style-type: none"><li>• Create, modify, or delete Panorama or firewall administrators and roles.</li><li>• Export, validate, revert, save, load, or import a configuration in the <b>Device &gt; Setup &gt; Operations</b> page.</li><li>• Configure <b>Scheduled Config Export</b> functionality in the <b>Panorama</b> tab.</li></ul>

**Admin role profiles:** To provide more granular access control over the functional areas of the web interface, CLI, and XML API, you can create custom roles. After new features are added to the product, you must update the roles with corresponding access privileges; the firewall and Panorama do not automatically add new features to the custom role definitions.

## 2.5.2 Different methods used to authenticate

### Authentication

Authentication is a method for protecting services and applications by verifying the identities of users so that only legitimate users have access. Several firewall and Panorama features require authentication. Administrators authenticate to access the web interface, CLI, or XML API of the firewall and Panorama. End users authenticate through Captive Portal or GlobalProtect to access various services and applications through the firewall. You can choose from several authentication services to protect your network and to accommodate the existing security infrastructure while ensuring a smooth user experience.

If you have a public key infrastructure (PKI), you can deploy certificates to enable authentication without requiring users to respond to login challenges manually. Alternatively, or in addition to certificates, you can implement interactive authentication, which requires users to authenticate using one or more methods.



Supported authentication types include the following:

- MFA
- SAML
- SSO
- Kerberos
- TACACS+
- RADIUS
- LDAP
- Local

### Protecting Service Access Through the Firewall

The Authentication policy enables you to authenticate end users before they can access services and applications. Whenever a user requests a service or application (such as by visiting a web page), the firewall evaluates the Authentication policy. Based on the matching Authentication policy rule, the firewall then prompts the user to authenticate using one or more methods (factors), such as login and password, voice, SMS, push, or OTP authentication. For the first factor, users authenticate through a Captive Portal web form. For any additional factors, users authenticate through an MFA login page.

After the user authenticates for all of the factors, the firewall evaluates the Security policy to determine whether or not to allow access to the service or application.

To reduce the frequency of authentication challenges that interrupt the user workflow, you can specify a timeout period during which a user authenticates only for initial access to services and applications, not for subsequent access. The Authentication policy integrates with Captive Portal to record the timestamps used to evaluate the timeout and to enable user-based policies and reports.

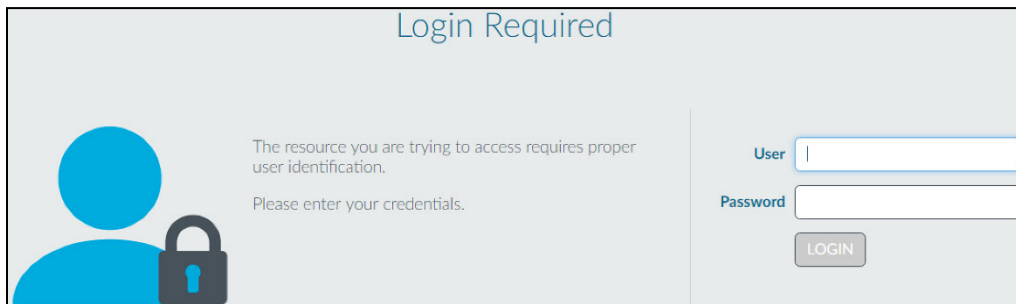
Based on the user information that the firewall collects during authentication, User-ID creates a new IP address-to-username mapping or updates the existing mapping for that user (if the mapping information has changed). The firewall generates User-ID logs to record the additions and updates. The firewall also generates an Authentication log for each request that matches an authentication rule. If you favor centralized monitoring, you can configure reports based on the User-ID or Authentication logs and forward the logs to Panorama or external services as you would for any other log types.

### Configuring the Authentication Policy

- Perform the following steps to configure Authentication policy for end users who access services through Captive Portal. Before starting, ensure that your Security policy allows users to access the services and URL categories that require authentication.
- Configure Captive Portal. If you use MFA services to authenticate users, you must set the **Mode** to **Redirect**.
- Configure the firewall to use one of the following services to authenticate users:
  - **External Authentication Services:** Configure a Server Profile to define how the firewall connects to the service.
  - **Local database authentication:** Add each user account to the local user database on the firewall.
  - **Kerberos SSO:** Create a Kerberos keytab for the firewall. You can configure the firewall to use Kerberos SSO as the primary authentication service and, if SSO failures occur, to fall back to an external service or local database authentication.

- Configure an Authentication Profile, an optional Authentication Sequence for each set of users, and Authentication policy rules that require the same authentication services and settings.
- Select the **Type** of authentication service and related settings:
  - **External service:** Select the **Type** of external server and select the **Server Profile** you created for it.
  - **Local database authentication:** Set the **Type** to **Local Database**. In the **Advanced settings**, **Add** the Captive Portal users and user groups you created.
  - **Kerberos SSO:** Specify the **Kerberos Realm** and **Import** the **Kerberos Keytab**.
- Configure an Authentication Enforcement object:
  - The object associates each Authentication Profile with a Captive Portal method. The method determines whether the first authentication challenge (factor) is transparent or requires a user response.
  - Select **Objects > Authentication** and **Add** an object.
  - Enter a **Name** to identify the object.
  - Select an **Authentication Method** for the authentication service type you specified in the Authentication Profile:
    - **browser-challenge:** Select this method if you want the client browser to respond transparently to the first authentication factor instead of the user entering login credentials. For this method, you must configure Kerberos SSO in the Authentication Profile or NTLM authentication in the Captive Portal settings. If the browser challenge fails, the firewall falls back to the web-form method.
    - **web-form:** Select this method if you want the firewall to display a Captive Portal web form for users to enter login credentials.
  - Select the Authentication Profile that you configured.
  - Enter the **Message** that the Captive Portal web form will display to tell users how to authenticate for the first authentication factor.
  - Click **OK** to save the object.
- Configure an Authentication policy rule:
  - Create a rule for each set of users, services, and URL categories that requires the same authentication services and settings.
  - Select **Policies > Authentication** and **Add** a rule.
  - Enter a **Name** to identify the rule.
  - Select **Source**, **Add** specific zones and IP addresses, or select **Any** zones or IP addresses.
  - Select **User** and select or **Add** the source users and user groups to which the rule applies (default is **any**).
  - Select or **Add** the **Host Information Profiles** to which the rule applies (default is **any**).
  - Select **Destination**, **Add** specific zones and IP addresses, or select any zones or IP addresses.
  - Select **Service/URL Category** and select or **Add** the services and service groups for which the rule controls access (default is **service-http**).

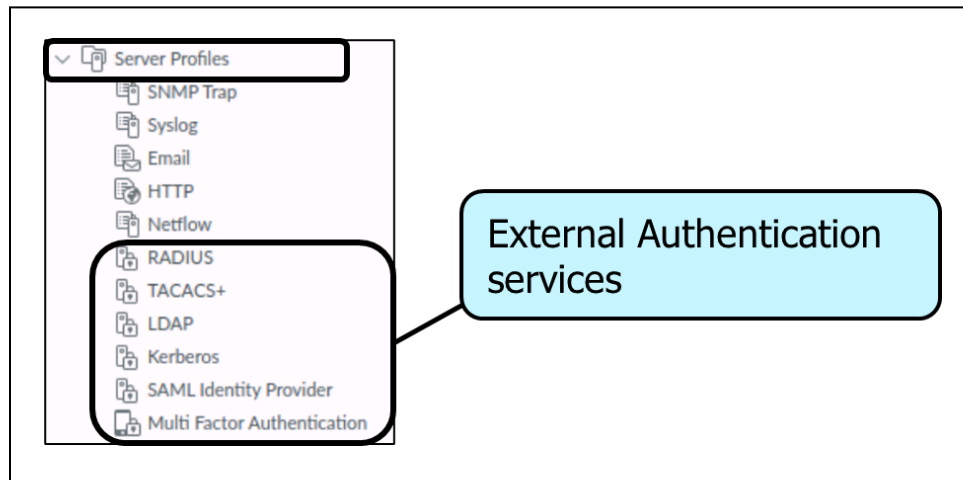
- Select or **Add** the **URL Categories** for which the rule controls access (default is **any**). For example, you can create a custom URL category that specifies the most-sensitive internal sites.
  - Select **Actions** and select the Authentication Enforcement object you created.
  - Specify the **Timeout** period in minutes (default is **60**) during which the firewall prompts the user to authenticate only once for repeated access to services and applications.
  - Click **OK** to save the rule.
- (MFA only) Customize the MFA login page:
    - The firewall displays this page so that users can authenticate for any additional MFA factors.
- Verify that the firewall enforces your Authentication policy:
    - Log in to your network as one of the source users specified in an Authentication policy rule.
    - Request a service or URL category that matches the one specified in the rule. The firewall displays the Captive Portal web form for the first authentication factor. Here is an example:



- End the session for the service or URL that you just accessed.
  - Start a new session for the same service or application. Be sure to perform this step within the timeout period that you configured in the Authentication rule.
  - The firewall allows access without re-authenticating.
  - Wait until the timeout period expires. Request the same service or application.
  - The firewall prompts you to re-authenticate.
- (Optional) Redistribute user mappings and authentication timestamps to the other firewalls that enforce the Authentication policy to ensure that they all apply timeouts consistently for all of the users.

### 2.5.3 The Authentication Sequence

When user or administrative access is configured, one or more authentication methods must be specified. A user or administrator definition typically requires an Authentication Profile that specifies the desired authentication method. When more than one method is desired, you can use an Authentication Sequence, which is a list of Authentication Profiles. The first profile will be accessed. If it is not available, the next option will be tried. An Authentication Profile specifies a single Server Profile. A Server Profile contains specific configuration and access information that is necessary to reach the external authentication service.



### 2.5.4 The device access method

#### Panorama Access Domains

Panorama access domains control the access that device group and template administrators have to specific device groups (to manage policies and objects), to templates (to manage network and device settings), and to the web interface of managed firewalls (through context switching). You can define up to 4,000 access domains, and you can manage them locally or by using the RADIUS Vendor-Specific Attributes (VSAs), TACACS+ VSAs, or SAML attributes.

### 2.5.5 References

- Configure an Authentication Profile and Sequence, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/authentication/configure-an-authentication-profile-and-sequence>
- Panorama > Access Domains, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/panorama-web-interface/panorama-access-domains>

## 2.6 Configure and manage certificates

### 2.6.1 Usage

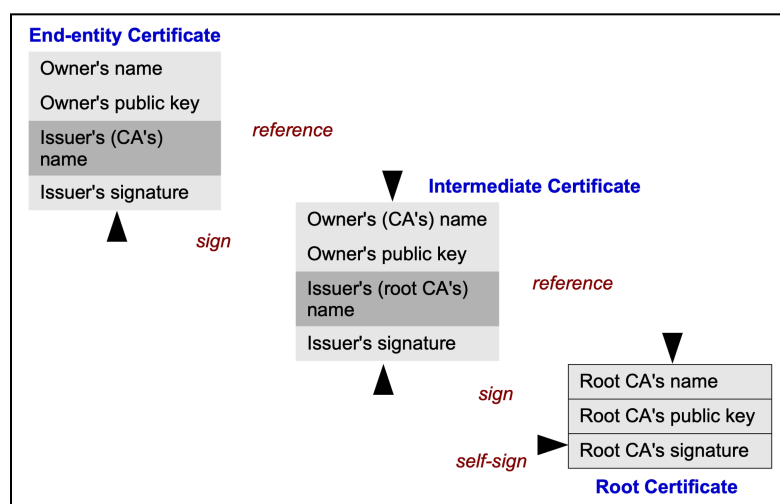
#### Certificate Background

In cryptography, a public key certificate, also known as a digital certificate or identity certificate, is an electronic document used to prove the ownership of a public key. The certificate includes information about the key, information about the identity of its owner (called the subject), and the digital signature of an entity that has verified the certificate's contents (called the issuer). If the signature is valid and the software examining the certificate trusts the issuer, then the software can use that key to communicate securely with the certificate's subject. In email encryption, code signing, and e-signature systems, a certificate's subject is typically a person or organization. However, in Transport Layer Security (TLS), a certificate's subject is typically a computer or other device, though TLS certificates may identify organizations or individuals in addition to their core role in identifying devices. TLS, sometimes called by its older name SSL, is notable for being a part of HTTPS, a protocol for securely browsing the web.

In a typical PKI scheme, the certificate issuer is a CA, usually a company that charges customers to issue certificates for them. CAs also can be created and managed by individuals and organizations requiring certificates for internal use.

A CA is responsible for signing certificates. These certificates act as an introduction between two parties, which means that a CA acts as a trusted third party. A CA processes requests from people or organizations requesting certificates (called subscribers), verifies the information, and potentially signs an end-entity certificate based on that information. To perform this role effectively, a CA needs to have one or more broadly trusted root certificates or intermediate certificates and the corresponding private keys. CAs may achieve this broad trust by having their root certificates included in popular software or by obtaining a cross-signature from another CA delegating trust.

A receiving entity is responsible for validating the information contained in a certificate presented to it. Among the potential verification tests is a validation that the certificate was actually issued by the CA whose information is in the certificate. This verification requires the CA's signing key contained in its Root Certificate used to sign all the issued certificates. This certificate must be locally available to the receiving entity to run the validation test. These CA Root Certificates are often kept in locally stored certificate caches in the hosting operating system or in a browser- or program-managed certificate cache. The firewall also contains a CA Root Certificate cache.



CAs also are responsible for maintaining up-to-date revocation information about the certificates they have issued, which indicates whether the certificates are still valid. They provide this information through Online Certificate Status Protocol (OCSP) or certificate revocation lists.

### 2.6.2 Profiles

Certificate profiles define user and device authentication for Authentication Portal, MFA, GlobalProtect, site-to-site IPsec VPN, external dynamic list validation, Dynamic DNS, User-ID agent and Terminal Services agent access, and web interface access to the Palo Alto Networks firewalls or Panorama. The profiles specify which certificates to use, how to verify the certificate revocation status, and how that status limits access. You need to configure a certificate profile for each application.

### 2.6.3 Chains

Not all of the websites send their complete certificate chain, even though the RFC 5246 TLSv1.2 standard requires authenticated servers to provide a valid certificate chain leading to an acceptable CA. When you enable decryption and apply a Forward Proxy Decryption Profile that enables block sessions with untrusted issuers in the decryption policy and if an intermediate certificate is missing from the certificate list the website's server presents to the firewall, the firewall can't construct the certificate chain to the top (root) certificate. In such cases, the firewall presents its forward untrust certificate to the client because the firewall cannot construct the chain to the root certificate and trust cannot be established without the missing intermediate certificate.

If a website you need to communicate with for business purposes has one or more missing intermediate certificates and the decryption policy blocks sessions with untrusted issuers, then you can find and download the missing intermediate certificate and install it on the firewall as a trusted root CA so that the firewall trusts the site's server. (The alternative is to contact the website owner and ask them to configure their server so that it sends the intermediate certificate during the handshake.)

### 2.6.4 References

Certificate Profiles,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/certificate-management/configure-a-certificate-profile>

Certificate Status,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/troubleshoot-and-monitor-decryption/decryption-logs/repair-incomplete-certificate-chains>

Keys and Certificates for Decryption Policies,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/decryption-concepts/keys-and-certificates-for-decryption-policies>

Certificate Management,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/certificate-management>

How to Install a Chained Certificate Signed by a Public CA,

<https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000ClkoCAC>

Resource List: SSL Certificates Configuring and Troubleshooting,


<https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000Cm5YCAS>

## 2.7 Configure routing

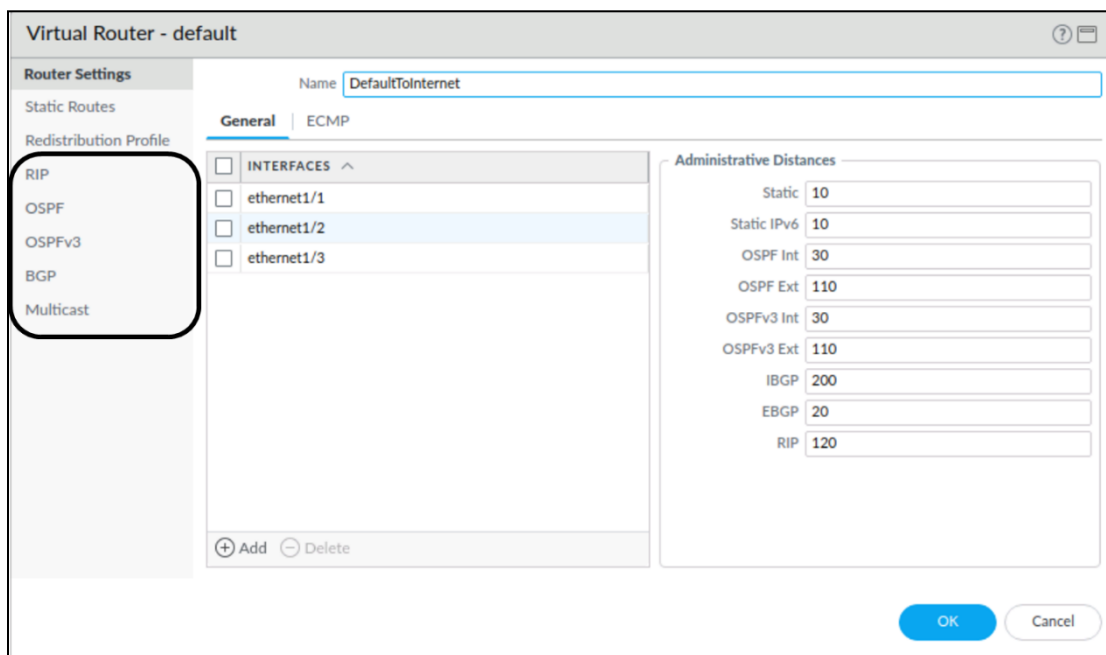
### 2.7.1 Dynamic routing

#### Traffic Forwarding


All of the traffic that arrives at the firewall is delivered either to an internal firewall process (destination traffic) or passed through a traffic interface (transit traffic). All of the transit traffic must be handed off to the egress interface by a traffic-handling object that matches the interface type. Examples of these objects are VLAN objects (VLANs) for Layer 2 traffic, virtual routers for Layer 3 traffic, and virtual wires for virtual wire interfaces.



Simultaneous implementations of multiple traffic handler types in multiple quantities are possible. Each object contains the configuration capabilities that are appropriate to its protocol-handling needs. Legacy virtual routers can implement various dynamic routing support, if desired. The Advanced Route Engine of virtual routers supports the Border Gateway Protocol (BGP) dynamic routing protocol and static routes.



Each Layer 3 dynamic routing protocol includes appropriate specific configuration options. Here is an example of the Open Shortest Path First (OSPF) protocol in the Legacy Route Engine:



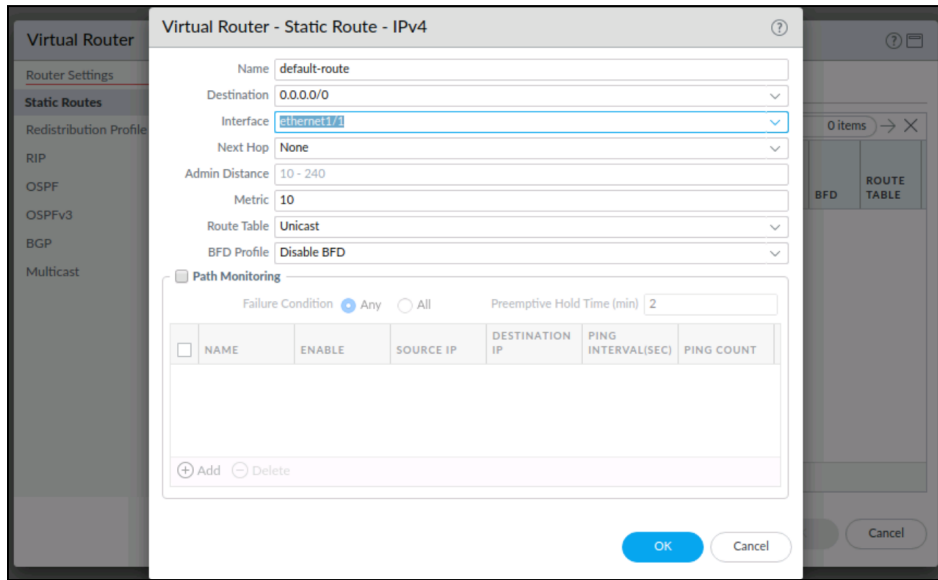
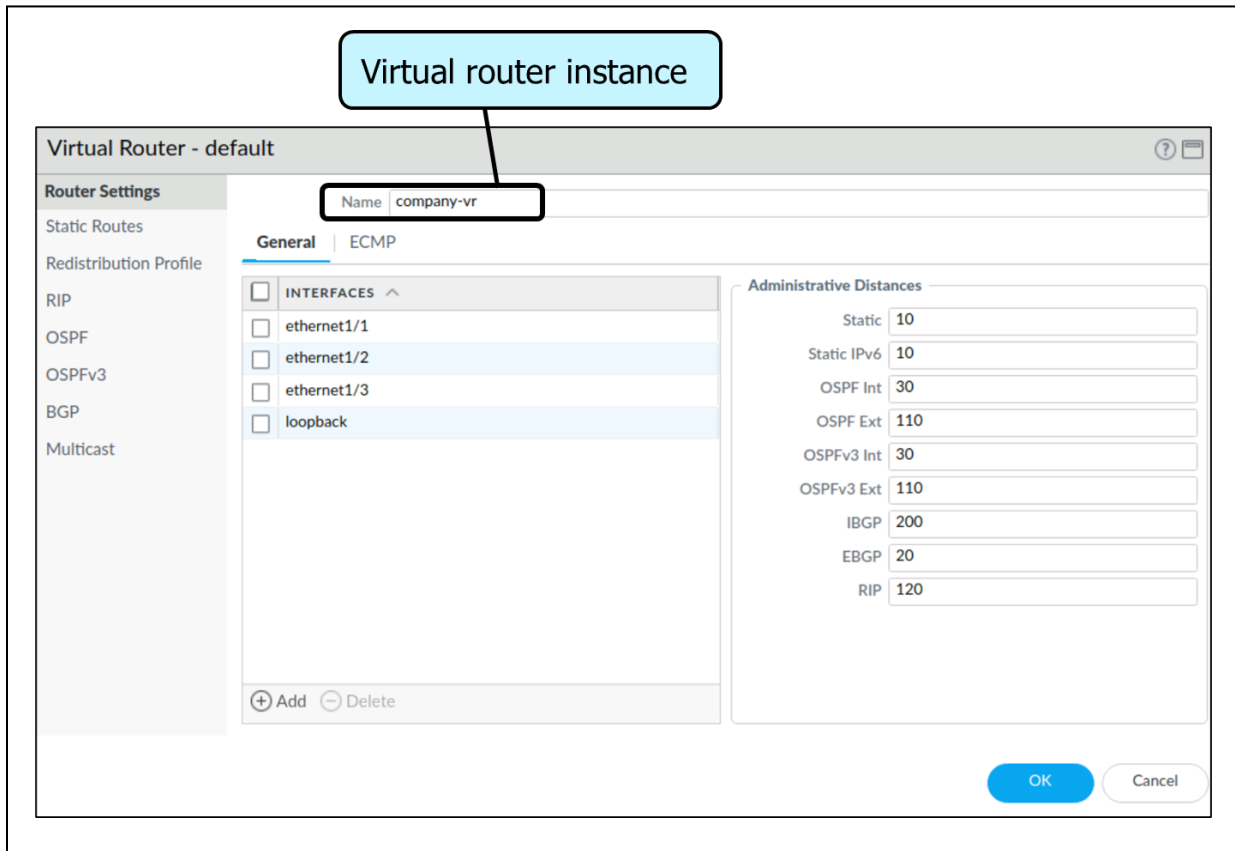
IPsec tunnels are considered Layer 3 traffic segments for implementation purposes and are handled by virtual routers like any other network segments. Forwarding decisions are made by destination address, not by VPN policy.


### Routing Configuration

PAN-OS supports static routes, BGP, OSPF, Routing Information Protocol (RIP), and multicast routing configured in two routing engines, only one of which can be active at a time. The Legacy Routing Engine is the continuation of the virtual routing features in the previous PAN-OS versions. It supports multiple dynamic routing protocols and can support more than one virtual routing instance, with the limit being determined by the firewall model. The Advanced Routing Engine only supports BGP and static routes and can support a single virtual router instance regardless of the firewall model. There are limitations for the number of entries in the forwarding tables (forwarding information bases [FIBs]) and the routing tables (routing information bases [RIBs]) in either routing engine.



The virtual router configuration is meant to match the existing routing infrastructure. In addition to protocol configuration, Redistribution Profiles can support protocol interoperability.





<input checked="" type="checkbox"/>	company-vr	ethernet1/1 ethernet1/2 ethernet1/3 loopback	Static Routes: 1 ECMP status: Disabled				<a href="#">More Runtime Stats</a>
-------------------------------------	------------	---	---	--	--	--	------------------------------------

## Virtual Routers

The firewall has two routing engines, one of which can be enabled at a time. The Legacy Route Engine is a continuation of the routing engine from the previous PAN-OS versions and is still the default. The Legacy Route Engine supports BGP, OSPF, OSPFv3, and RIP dynamic routing protocols, plus static routes, route monitoring, and Redistribution Profiles. Several virtual router instances can be created and managed simultaneously. The Advanced Route Engine is also available in some firewall models and supports the BGP dynamic routing protocol only with static routes. The Advanced Route Engine allows for only a single virtual router instance. A firewall must be rebooted when the type of route engine is changed. Firewalls that use the Advanced Route Engine are appropriate for large data centers, enterprises, ISPs, and cloud services.

## Administrative Distance

Within the virtual router configuration, set administrative distances for types of routes as required for your network. A virtual router that has two or more different routes to the same destination uses an administrative distance to choose the best path from different routing protocols and static routes by preferring a lower distance.

## ECMP Routing

ECMP processing is a networking feature that enables the firewall to use up to four equal-cost routes to the same destination. Without this feature, the virtual router selects only a single route to a destination from the routing table and adds it to its forwarding table; it does not use any of the other routes unless there is an outage in the chosen route.

Enablement of ECMP functionality on a virtual router allows the firewall to have up to four equal-cost paths to a destination in its forwarding table, which allows the firewall to perform these actions:

- Load balance flows (sessions) to the same destination over multiple equal-cost links
- Efficient use of all the available bandwidth on the links to the same destination rather than leaving some links unused
- Dynamic shifting of traffic to another ECMP member to the same destination if a link fails rather than waiting for the routing protocol or RIB table to elect an alternative path/route, which helps reduce downtime

## 2.7.2 Redistribution Profiles

### Route Redistribution

Route redistribution on the firewall is the process of making routes that the firewall learned from one routing protocol (or a static or connected route) available to a different routing protocol, thereby increasing the number of reachable networks. Without route redistribution, a router or virtual router advertises and shares routes only with the other routers that run the same routing protocol. You can redistribute IPv4 or IPv6 BGP, connected, or static routes into the OSPF RIB and redistribute OSPFv3, connected, or static routes into the BGP RIB.

Route distribution means that, for example, you can make specific networks which were earlier available only by manual static route configuration on specific routers now available to BGP autonomous systems or OSPF areas. You can also advertise locally connected routes—such as routes to a private lab network—into the BGP autonomous systems or OSPF areas.

You might want to give users on your internal OSPFv3 network access to BGP so they can access devices on the internet. In this case, you would redistribute BGP routes into the OSPFv3 RIB.

Conversely, you might want to give your external users access to some parts of your internal network, so you can make internal OSPFv3 networks available through BGP by redistributing OSPFv3 routes into the BGP RIB.

## 2.7.3 Static routes

When you configure static routes, they are normally used with dynamic routing protocols. Typically, you configure a static route for a location that a dynamic routing protocol can't reach.

## 2.7.4 Route monitoring

When you configure path monitoring for a static route, the firewall uses path monitoring to detect when the path to the monitored destination has gone down. The firewall then reroutes traffic by using alternative routes.

## 2.7.5 Policy-based forwarding

The firewall in most cases uses the destination IP address in a packet to determine the egress interface. The firewall uses the routing table associated with the virtual router to which the interface is connected to perform the route lookup. Policy-based forwarding (PBF) allows you to override the routing table. You can specify the egress interface-based set parameters (such as destination IP address) or type of traffic.

When you create a PBF rule, you must specify:

- A name for the rule
- A source zone or interface
- An egress interface

You can specify the source and destination addresses by using an IP address, an address object, or a FQDN. Note that application-specific rules are not recommended for use with PBF because PBF rules might be applied before the firewall has determined the application.

## 2.7.6 Virtual routers versus logical routers

### Virtual Routers

Layer 3 interfaces and their associated virtual routers are the most widely used deployment options.

A virtual router is a function of the firewall that participates in Layer 3 routing. The firewall uses virtual routers to obtain routes to other subnets after you manually define static routes or through participation in one or more Layer 3 routing protocols (dynamic routes). The routes that the firewall obtains through these methods populate the IP RIB on the firewall. When a packet is destined for a different subnet than the one it arrived on, the virtual router obtains the best route from the RIB, places it in the FIB, and forwards the packet to the next hop router that is defined in the FIB. The firewall uses Ethernet switching to reach other devices on the same IP subnet. (An exception to adding only a single optimal route to the FIB occurs if you are using ECMP, in which case all equal-cost routes go in the FIB.)

The Ethernet, VLAN, and tunnel interfaces that are defined on the firewall receive and forward Layer 3 packets. The destination zone is derived from the outgoing interface based on the forwarding criteria, and the firewall consults policy rules to identify the Security policies that it applies to each packet. 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.

You can configure Layer 3 interfaces on a virtual router to participate with dynamic routing protocols (BGP, OSPFv2, OSPFv3, or RIP) and add static routes with the routing protocol configured in the routing engine. You can also create multiple virtual routers in the Legacy Route Engine; each router maintains a separate set of routes that are not shared by the other virtual routers, which enables you to configure different routing behaviors for different interfaces.

Each Layer 3 Ethernet, loopback, VLAN, and tunnel interface defined on the firewall must be associated with a virtual router. Although each interface can belong to only one virtual router, you can configure multiple routing protocols and static routes for a virtual router.

A firewall can have more than one router instance when it is using the Legacy Route Engine, with each model supporting a different maximum. An interface can be attached to one virtual router at a time. Virtual routers can route directly to each other within the firewall.

### **Logical Routers**

The firewall uses logical routers to obtain Layer 3 routes to other subnets when you manually define static routes or through participation in one or more Layer 3 routing protocols (dynamic routes). The routes that the firewall obtains through these methods populate the IP RIB on the firewall. When a packet is destined for a different subnet than the one it arrived on, the logical router obtains the best route from the RIB, places it in the FIB, and forwards the packet to the next hop router defined in the FIB. The firewall uses Ethernet switching to reach other devices on the same IP subnet. (An exception to one best route going in the FIB occurs if you are using ECMP, in which case all the equal-cost routes go in the FIB.)

The Ethernet, VLAN, and tunnel interfaces defined on the firewall receive and forward Layer 3 packets. The destination zone is derived from the outgoing interface based on the forwarding criteria, and the firewall consults policy rules to identify the Security policies that it applies to each packet. In addition to routing to other network devices, logical routers can route to other logical routers within the same firewall if a next hop is specified to point to another logical router.

You can configure Layer 3 interfaces to participate with the dynamic routing protocols (BGP, OSPF, OSPFv3, or RIP) as well as to add static routes. You can also create multiple logical routers, each maintaining a separate set of routes that aren't shared by logical routers, enabling you to configure different routing behaviors for different interfaces.

You can configure dynamic routing from one logical router to another by configuring a loopback interface in each logical router, creating a static route between the two loopback interfaces, and then configuring a dynamic routing protocol to peer between these two interfaces.

Each Layer 3 Ethernet, loopback, VLAN, and tunnel interface defined on the firewall must be associated with a logical router. While each interface can belong to only one logical router, you can configure multiple routing protocols and static routes for a logical router. Regardless of the static routes and dynamic routing protocols you configure for a logical router, one general configuration is required.

## 2.7.7 References

- Configure a Static Route,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/static-routes/configure-a-static-route>
- Static Route Removal Based on Path Monitoring,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/static-routes/static-route-removal-based-on-path-monitoring>
- Service Versus Applications in PBF,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/policy/policy-based-forwarding/pbf/service-versus-applications-in-pbf>
- How to Configure PBF in Multi Vsys Configuration,  
<https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIKsCAK>
- Network > Virtual Routers,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/network/network-virtual-routers>
- ECMP,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/ecmp>

## 2.8 Configure NAT

### 2.8.1 NAT policy rules

NAT allows the organization to use internal IP addresses that are not exposed to the Internet. NAT rules are based on source and destination zones, source and destination addresses, and application services (such as HTTP). As with Security policy rules, NAT policy rules are compared against incoming traffic in sequence and the first rule that matches the traffic is applied.

### 2.8.2 Security rules

A Security policy allows you to enforce rules and actions. It can be as general or as 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 the applications if all of the other traffic-related settings are the same.

Security policy rules are matched from top down. Up to two processing steps are in each Security policy match. Step 1 confirms that a match has been made based on the matching conditions provided in the Security policy. If a match is found in Step 1, the traffic is logged (based on that policy rule's configuration) and the chosen action (deny, allow, drop, or reset) is performed. Once processing is complete, no further matching is done in the Security policy rulebase.

### 2.8.3 Source NAT

Source NAT is typically used by internal users to access the internet; the source address is translated and kept private. There are three types of source NAT: dynamic IP and port (DIPP), dynamic IP, and static IP.

## DIPP

DIPP allows multiple hosts to have their source IP addresses translated to the same public IP address with different port numbers. The dynamic translation is to the next available address in the NAT address pool, which you configure as a translated address pool be to an IP address, range of addresses, a subnet, or a combination of these. As an alternative to using the next address in the NAT address pool, DIPP allows you to specify the address of the interface itself. The advantage of specifying the interface in the NAT rule is that the NAT rule is automatically updated to use any address subsequently acquired by the interface.

DIPP is sometimes referred to as interface-based NAT or network address port translation. DIPP has a default NAT oversubscription rate, which is the number of times that the same translated IP address and port pair can be used concurrently.

## Dynamic IP

Dynamic IP allows the one-to-one, dynamic translation of a source IP address only (no port number) to the next available address in the NAT address pool. The size of the NAT pool should be equal to the number of internal hosts that require address translations. By default, if the source address pool is larger than the NAT address pool and eventually all of the NAT addresses are allocated, new connections that need address translation are dropped. To override this default behavior, use Advanced (Dynamic IP/Port Fallback) to enable use of DIPP addresses when necessary. In either event, as sessions terminate and the addresses in the pool become available, they can be allocated to translate new connections. Dynamic IP NAT supports the option for you to reserve dynamic IP NAT addresses.

## Static IP

Static IP allows the one-to-one, static translation of a source IP address, but leaves the source port unchanged. A common scenario for a static IP translation is an internal server that must be available to the internet.

### 2.8.4 No-NAT Policies

No-NAT rules are configured to allow the 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.

You can verify the NAT rules processed by selecting **Device > Troubleshooting** and testing the traffic matches for the NAT rule. For example:

Test Configuration	Test Result	Result Detail				
<p>Select Test: NAT Policy Match</p> <p>From: I3-vlan-trust</p> <p>To: I3-untrust</p> <p>Source: 10.54.21.28</p> <p>Destination: 8.8.8.8</p> <p>Source Port: [ 1 - 65535 ]</p> <p>Destination Port: 445</p> <p>Protocol: 6</p> <p>To Interface: None</p> <p>Ha Device ID: [ 0 - 1 ]</p> <p><input type="button" value="Execute"/> <input type="button" value="Reset"/></p>	<p>NAT Policy Match Result</p>	<table border="1"><thead><tr><th>Name</th><th>Value</th></tr></thead><tbody><tr><td>Result</td><td>access-corp</td></tr></tbody></table>	Name	Value	Result	access-corp
Name	Value					
Result	access-corp					

### 2.8.5 Use session browser to find NAT rule name

The session browser allows you to browse and filter current running sessions on the firewall. In the details of a session, you can see which NAT Policy rule is being used to process a session.

### 2.8.6 U-Turn NAT

The term “U-Turn” is used when the logical path of a connection traverses the firewall from inside to outside and back in by connecting to an internal resource by using its external IP address. U-Turn NAT is a configuration trick to accommodate a deployment where the external IP needs to reach an internal resource.

#### Use of U-turn NAT

In some environments, an internal host might require an external IP address to run a certain service—for example, a locally hosted web server or mail server. Internal hosts may need to use the external IP address due to the absence of an internal DNS server or other requirements specific to the service.

In this example, using regular destination NAT configuration, any connections originating from the laptop directed to the server on its external IP address, 198.51.100.230, are directed to the default gateway because the IP address is not in the local subnet. Connections then get translated to the destination IP address 192.168.0.97 without applying the source NAT, which causes the web server to send return packets directly to the workstation, resulting in an asymmetric flow.

With U-Turn NAT configured, outbound packets from the laptop also have the source NAT applied to them. The source NAT causes the server to send reply packets directly to the firewall rather than to the laptop. Sending packets directly to the firewall prevents asymmetry and allows the firewall to still apply content scanning to the session.

#### Configuring U-Turn NAT

The Security policy has an inbound rule that allows inbound connections from the internet onto the internal web server with application web browsing, which is default port 80. Further, the outbound Security policy allows all the users to go to the internet on any application. In addition, two implied rules allow intrazonal traffic—for example, trust to trust—and the denied intranet zone prevents sessions from reaching other zones without an explicit policy permitting it.

The NAT policy has an inbound rule to allow connections from anywhere to the external IP address to be translated to the server’s internal IP address. It also has a hide-NAT rule to allow internal connections to go out to the internet and get source-translated behind the firewall’s external interface IP address.

When a client PC is trying to access 198.51.100.230 (the internet-facing IP address of the internal server), it will result in a page not loading. Wireshark shows a syn packet being sent to the external IP, a syn/ack being received from the internal IP address 192.168.0.97, and a reset being sent because the client doesn’t understand what’s going on.

If we now go back to the firewall and open the NAT policy, we see that the inbound NAT rule has been set to accept any source zone and translate that to the proper internal server IP address.



## Creating a New NAT Rule Details:

- **Name:** internal access
- **Source zone:** trust
- **Destination zone:** untrust
- **Destination address:** 198.51.100.230

## Under the Translated Packet tab:

- **Destination address:** 192.168.0.97 (IP address of the web server in question)
- **Source address translation:** Dynamic IP/Port
- **Switch address type:** Interface
- **Interface:** ethernet1/2 (internal interface of the firewall)
- **IP address:** 192.168.0.230/24

Name this new rule internal access. Go to the **Original Packet** tab and set the **Source Zone** to **trust**, the **Destination Zone** to **untrust**, and the **Destination Address** to **198.51.100.230**. In the **Translation Packet** tab, set the **Destination Address**—just like the regular rule—to **192.168.0.97**. Enable source address translation by setting it to **Dynamic IP and Port** and switch the **Address Type** to **interface address**.

You also can set the **Address Type** to **translated address** and choose an address in the IP range assigned to the interface. In this example, let's continue with the IP address assigned to the interface for ease of use.

Select the trust zone interface from the drop-down list, set its IP, and click **OK**.

**NOTE:** Be sure to place the new NAT rule above the inbound rule. Otherwise, the original NAT rule will take precedence over the newly created rule.

Commit the configuration and return to the client PC.

## Verifying and Testing U-Turn NAT

If we open the web page now, the internet information server 7 default page loads and the web server is accessible from the inside on its external IP address.

If we take a look at the Wireshark packet capture, the client is receiving its returning packets from the external IP because the firewall can now perform NAT on both directions of the flow.

### 2.8.7 Check HIT counts

View the number of times a Security, NAT, QoS, policy-based forwarding (PBF), Decryption, Tunnel Inspection, Application Override, Authentication, or DoS protection rule matches traffic to help keep firewall policies up to date as the environment and security needs change. To prevent attackers from exploiting over-provisioned access, such as when a server is decommissioned or when you no longer need temporary access to a service, use the policy rule hit count data to identify and remove unused rules.

Policy rule usage data enables you to validate rule additions and rule changes and monitor the time frame when a rule was used. For example, when you migrate port-based rules to app-based rules, you create an app-based rule above the port-based rule and check for any traffic that matches the port-based rule. After migration, the hit count data helps you determine whether it is safe to remove the port-based rule by confirming whether traffic is matching the app-based rule instead of the port-based rule. The policy rule hit count helps you determine whether a rule is effective for access enforcement.

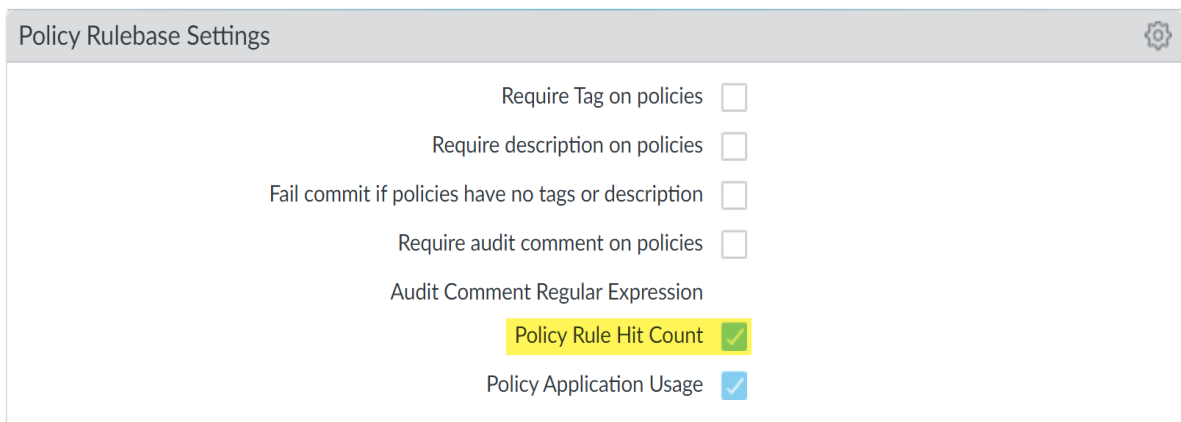
You can reset the rule hit count data to validate an existing rule or to gauge rule usage within a specified period of time. Policy rule hit count data is not stored on the firewall or Panorama so that data is no longer available after you reset (clear) the hit count.

After filtering the policy rulebase, administrators can take action to delete, disable, enable, and tag policy rules directly from the policy optimizer. For example, you can filter unused rules and then tag them for review to determine whether they can be safely deleted or kept in the rulebase. By enabling administrators to take action directly from the policy optimizer, you can reduce the management overhead required to further assist in simplifying rule lifecycle management and ensure that firewalls are not over-provisioned.

Perform the following high-level tasks if required:

**Step 1:** Launch the web interface.

**Step 2:** Verify that **Policy Rule Hit Count** is enabled.



**Step 3:** Select **Policies**.

**Step 4:** View the policy rule usage for each policy rule:

- **Hit Count** — The number of times traffic matched the criteria defined in the policy rule and persists through reboot, data plane restarts, and upgrades unless you manually reset or rename the rule.
- **Last Hit** — The most recent timestamp for when traffic matched the rule.
- **First Hit** — The first instance when traffic was matched to this rule.
- **Modified** — The date and time the policy rule was last modified.
- **Created** — The date and time the policy rule was created.

NAME	Source				Rule Usage			MODIFIED	CREATED
	T...	Z...	A...	U...	HIT COUNT	LAST HIT	FIRST HIT		
Video	n...	a...	a...	a...	2424328	2020-09-22 11:33:00	2019-07-30 10:12:57	2020-07-27 13:27:16	2019-07-30 09:50
Video Streaming	n...	a...	a...	a...	14337228	2020-09-22 16:26:58	2019-07-30 10:12:57	2020-07-27 13:27:16	2019-07-30 09:50
Wavenger	n...	a...	a...	a...	321760616	2020-09-22 16:27:10	2019-07-30 10:12:57	2020-07-27 13:27:16	2019-07-30 09:50
Web Traffic	n...	a...	a...	a...	1509584361	2020-09-22 16:27:10	2019-07-30 10:12:02	2020-07-27 13:27:16	2019-07-30 09:50
iperf	n...	a...	a...	a...	5	2019-10-15 14:54:31	2019-10-11 13:08:28	2020-07-27 13:27:16	2019-07-30 09:50

**Step 5:** In the Policy Optimizer dialog, view the **Rule Usage** filter.

**Step 6:** Filter rules in the selected rulebase.

1. Select the **Timeframe** you want to filter on or specify a **Custom** time frame.
2. Select the rule **Usage** on which to filter.
3. (Optional) If you have reset the rule usage data for any rules, check for **Exclude rules reset during the last <number of days> days** and decide when to exclude a rule based on the number of days you specify since the rule was reset. Only rules that were reset before the specified number of days are included in the filtered results.

NAME	HIT COUNT	LAST HIT	FIRST HIT	RESET DATE	MODIFIED	CREATED
1 Deny_Malicious	75211831	2020-06-24 10:58:26	2019-08-13 14:38:29	-	2020-07-27 13:27:16	2019-07-30 09:50:23
2 Block_Quic	2809657	2020-09-11 00:15:57	2019-08-22 08:14:02	-	2020-07-27 13:27:16	2019-07-30 09:50:23
3 Allow_DNS	433179426	2020-09-22 16:35:47	2019-08-13 14:39:37	-	2020-07-27 13:27:16	2019-07-30 09:50:23
4 Block_PasteBin_Reddit...	18290041	2020-09-22 16:33:45	2020-04-15 18:00:36	-	2020-07-27 13:27:16	2020-04-15 17:29:12
5 Block_Social Media	0	-	-	-	2020-07-27 13:27:16	2020-06-30 16:37:15
6 Temp Allow for Cont...	0	-	-	-	2020-07-27 13:27:16	2020-05-22 17:35:44
7 Allow_Fetch	161307	2020-08-13 09:34:46	2020-04-15 18:45:07	-	2020-07-27 13:27:16	2020-04-15 18:44:46
8 Allow_SCADA_Traffic	357362	2020-09-22 16:35:09	2020-04-09 11:34:44	-	2020-07-27 13:27:16	2020-04-09 11:34:48
9 Zoom	0	-	-	-	2020-07-27 13:27:16	2020-04-16 11:43:49
10 Allow_Goalie	4974276	2020-09-22 16:18:20	2020-04-16 11:48:02	-	2020-07-27 13:27:16	2020-04-16 11:43:49
11 Allow_Office365_Core	235	2020-09-22 13:19:47	2020-05-22 17:49:50	-	2020-07-27 13:27:16	2020-05-22 17:28:26
12 Allow_Office365_Infra	0	-	-	-	2020-07-27 13:27:16	2020-05-22 22:46:44
13 Allow_Office365_ssf...	29597	2020-09-22 16:33:01	2020-05-22 22:55:02	-	2020-07-27 13:27:16	2020-05-22 22:46:44
14 Allow_March_Madness	13980	2020-08-11 08:54:17	2020-04-09 15:22:46	-	2020-07-27 13:27:16	2020-04-09 16:47:09
15 Allow_ssl_Http	33526900	2020-09-22 16:33:45	2020-04-09 15:22:46	-	2020-07-27 13:27:16	2020-04-09 16:47:09
16 Known_Device_Ping	151859	2020-08-13 09:36:37	2020-04-13 16:57:45	-	2020-07-27 13:27:16	2020-04-13 16:39:40
17 Allow_Office_Interna...	30	2020-08-13 09:36:56	2020-04-22 11:26:54	-	2020-07-27 13:27:16	2020-04-22 11:26:20

4. (Optional) Specify search filters based on rule data.

- Hover your cursor over the column header and **Columns**.
- Add any additional columns you want to display or use for the filter.

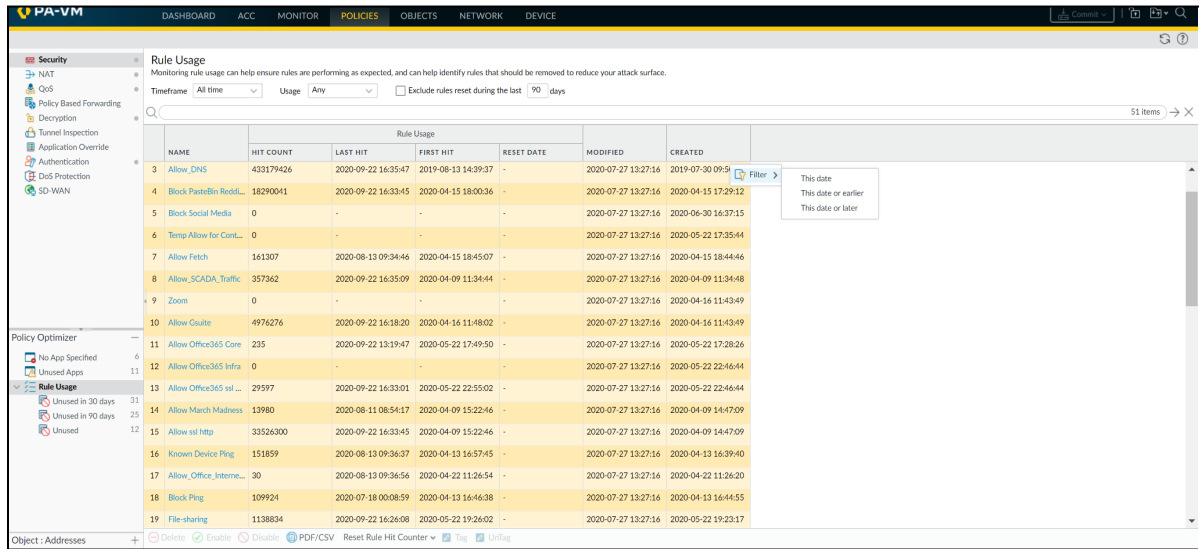
The screenshot shows a table with a 'CREATED' column header. A dropdown menu is open over the 'Columns' header, displaying a list of columns with checkboxes. The table contains several rows of data, each with a timestamp in the 'CREATED' column.

CREATED	
2020-04-15 17:28:	
2020-06-03 16:02:	
2020-05-22 17:34:57	
2020-04-15 18:43:40	
2020-04-09 11:34:03	
2020-04-16 11:42:46	
2020-04-16 11:42:46	
2020-05-22 17:26:44	
2020-05-22 22:45:53	
2020-05-22 22:45:53	
2020-04-09 14:44:37	
2020-04-09 14:44:37	
2020-04-13 16:38:36	
2020-04-22 11:25:01	
2020-04-13 16:43:49	

Columns menu items:

- Name
- Location
- Service
- Tags
- Type
- Source Zone
- Source Address
- Source User
- Source
- Destination Zone
- Destination Address
- Application
- URL Category
- Action
- Profile
- Options
- Rule UUID
- Target
- Description
- Traffic (Bytes, 30 days)
- App Usage Apps Allowed
- App Usage Apps Seen
- App Usage Days with No New Apps
- App Usage Compare
- Rule Usage
- Modified
- Created

- Hover your cursor over the column data that you would like to filter on **Filter**. For data that contain dates, select whether to filter by using **This date**, **This date or earlier**, or **This date or later**.
- **Apply Filter** ( → ).



### Step 7: Take action on one or more unused policy rules.

1. Select one or more unused policy rules.
2. Perform one of the following actions:
  - a. **Delete** — Delete one or more selected policy rules.
  - b. **Enable** — Enable one or more selected policy rules when disabled.
  - c. **Disable** — Disable one or more selected policy rules.
  - d. **Tag** — Apply one or more group tags to one or more selected policy rules. The group tag must already exist in order to tag the policy rule.
  - e. **Untag** — Remove one or more group tags from one or more selected policy rules.
3. Commit your changes.

## 2.8.7 References

- View Policy Rule Usage, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/view-policy-rule-usage>
- Palo Alto Networks #1: Initial Configuration (for beginners), <https://rtodto.net/palo-alto-networks-1-initial-configuration/>
- NAT Policy Overview, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/nat/nat-policy-rules/nat-policy-overview>
- Configure NAT, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/nat/configure-nat>
- NAT Configuration Examples, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/nat/nat-configuration-examples>
- Policies > NAT, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/policies/policies-nat>
- Configure Session Settings, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/session-settings-and-timeouts/configure-session-settings>

## 2.9 Configure site-to-site tunnels

To set up a VPN tunnel, the Layer 3 interface at each end must have a logical tunnel interface for the firewall to connect to and establish a VPN tunnel. A tunnel interface is a logical (virtual) interface that is used to deliver traffic between two endpoints. If you configure any proxy IDs, the proxy ID is counted toward any IPSec tunnel capacity.

The tunnel interface must belong to a security zone to apply policy, and it must be assigned to a virtual router in order to use the existing routing infrastructure. Ensure that the tunnel interface and the physical interface are assigned to the same virtual router so that the firewall can perform a route lookup and determine the appropriate tunnel to use.

Typically, the Layer 3 interface that the tunnel interface is attached to belongs to an external zone, such as the untrust zone. While the tunnel interface can be in the same security zone as the physical interface, you can create a separate zone for the tunnel interface for added security and better visibility. If you create a separate zone for the tunnel interface—for example, VPN zone—you will need to create security policies to enable traffic to flow between the VPN zone and the trust zone.

To route traffic between the sites, a tunnel interface does not require an IP address. An IP address is only required if you want to enable tunnel monitoring or if you are using a dynamic routing protocol to route traffic across the tunnel. With dynamic routing, the tunnel IP address serves as the next hop IP address for routing traffic to the VPN tunnel.

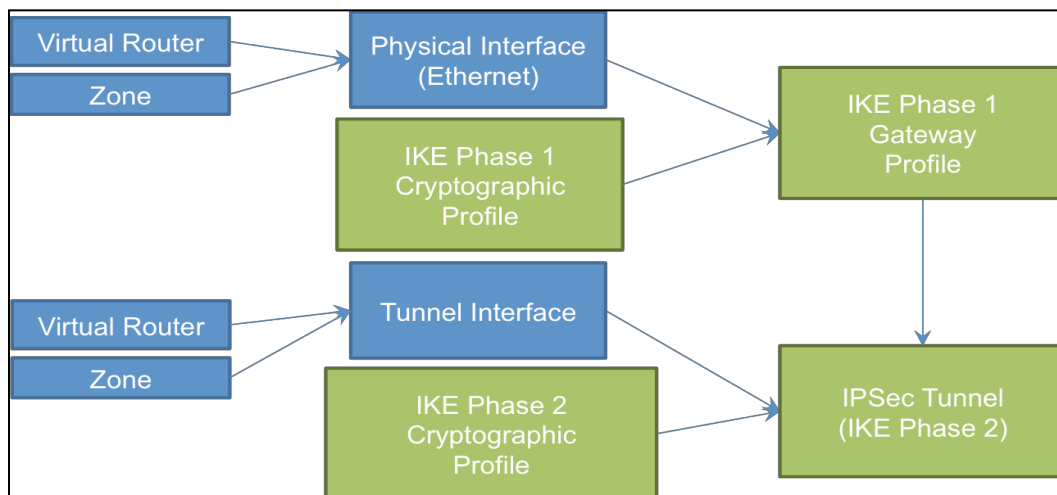
If you are configuring the Palo Alto Networks firewall with a VPN peer that performs policy-based VPN, you must configure a local and remote Proxy ID when setting up the IPSec tunnel. Each peer

compares the Proxy-IDs configured on it with what is actually received in the packet to allow a successful IKE phase 2 negotiation. If multiple tunnels are required, configure unique Proxy IDs for each tunnel interface; a tunnel interface can have a maximum of 250 Proxy IDs. Each Proxy ID counts toward the IPsec VPN tunnel capacity of the firewall, and the tunnel capacity varies according to the firewall model.

### 2.9.1 IPsec components

#### IPsec Tunnel Interfaces

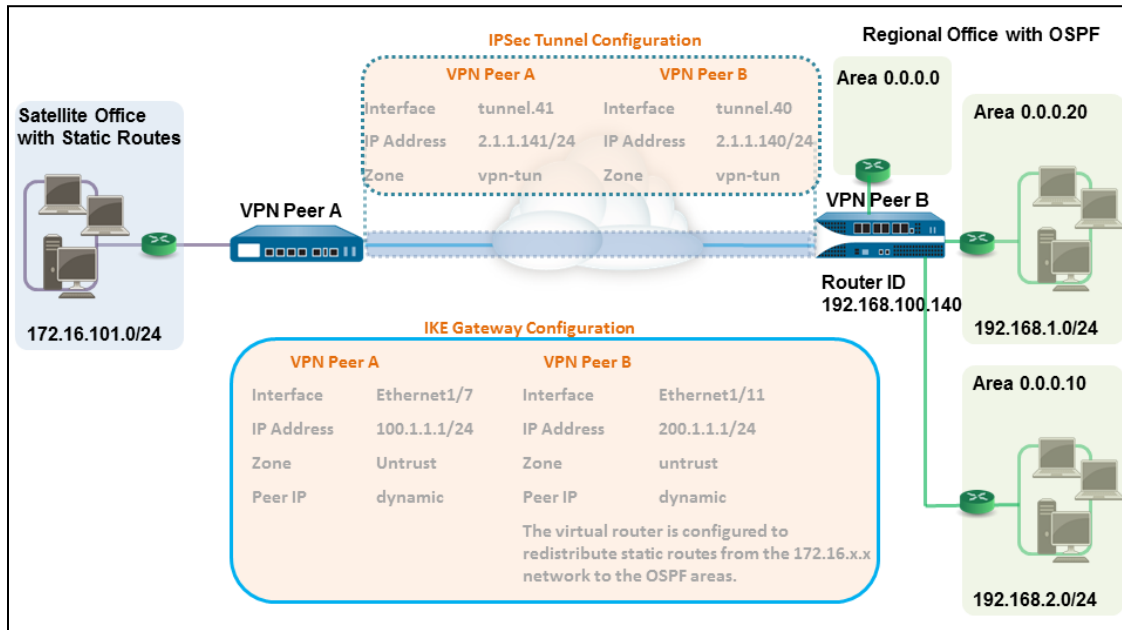
IPsec VPNs are terminated on Layer 3 tunnel interfaces. (These tunnel interfaces can be put into separate zones, thus allowing a specific Security policy per zone.) These tunnels require IPsec and Crypto Profiles for Phase 1 and Phase 2 connectivity. PAN-OS supports route-based VPNs, which means that the decision to route traffic through the VPN is made by the virtual router. Palo Alto Networks firewalls support connections to alternative policy-based VPNs, thus requiring the use of proxy IDs for compatibility. The following figure shows the various objects involved in IPsec tunnel definitions:



### 2.9.2 Static peers and dynamic peers for IPsec

In this example, one site uses static routes and the other site uses OSPF. When the routing protocol is not the same between locations, the tunnel interface on each firewall must be configured with a static IP address. Then, to allow the exchange of routing information, the firewall that participates in both the static and dynamic routing processes must be configured with a Redistribution Profile. Configuring the Redistribution Profile enables the virtual router to redistribute and filter routes between protocols—static routes, connected routes, and hosts—from the static autonomous system to the OSPF autonomous system. Without this Redistribution Profile, each protocol functions on its own and does not exchange any route information with other protocols running on the same virtual router.

In this example, the satellite office has static routes and all the traffic destined to the 192.168.x.x network is routed to tunnel.41. The virtual router on VPN Peer B participates in both static and dynamic routing processes and is configured with a Redistribution Profile to propagate (export) the static routes to the OSPF autonomous system.



### 2.9.3 IPsec tunnel Monitor Profiles

A Monitor Profile is used to monitor IPsec tunnels and a next-hop device for 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 they can be very useful for maintaining connectivity between sites and ensuring that PBF rules are maintained. The following settings are used to configure a Monitor Profile.

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. <ul style="list-style-type: none"> <li><b>wait-recover:</b> Wait for the tunnel to recover; do not take additional action. Packets will continue to be sent according to the PBF rule.</li> <li><b>fail-over:</b> Traffic will fail over to a backup path, if one is available. The firewall uses the routing table lookup to determine routing for the duration of this session.</li> </ul> In both cases, the firewall tries to negotiate new IPsec keys to accelerate recovery.
Interval	Specify the time between heartbeats (range is 2 to 10; default is 3).
Threshold	Specify the number of heartbeats to be lost before the firewall takes the specified action (range is 2 to 10; default is 5).



## 2.9.4 IPsec tunnel testing

Perform this task to test VPN connectivity.

**Step 1:** Initiate IKE phase 1 by either pinging a host across the tunnel or using the following CLI command:

```
test vpn ike-sa gateway <gateway_name>
```

**Step 2:** Enter the following command to test if IKE phase 1 is set up:

```
show vpn ike-sa gateway <gateway_name>
```

In the output, check whether the Security Association displays. If it doesn't, review the syslog messages to interpret the reason for failure.

**Step 3:** Initiate IKE phase 2 by either pinging a host from across the tunnel or using the following CLI command:

```
test vpn ipsec-sa tunnel <tunnel_name>
```

**Step 4:** Enter the following command to test if IKE phase 2 is set up:

```
show vpn ipsec-sa tunnel <tunnel_name>
```

In the output, check whether the Security Association displays. If it doesn't, review the syslog messages to interpret the reason for failure.

**Step 5:** To view the VPN traffic flow information, use the following command:

```
show vpn flow
total tunnels configured:          1
filter - type IPSec, state any

total IPSec tunnel configured:    1
total IPSec tunnel shown:        1

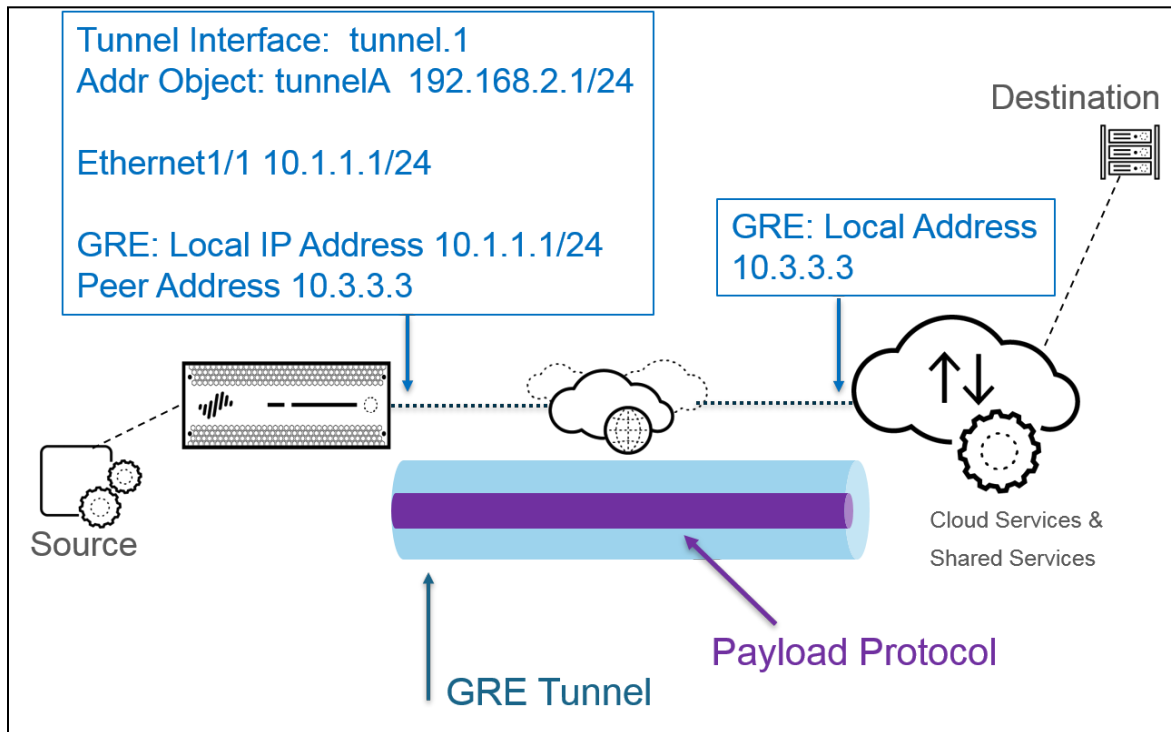
name          id    state  local-ip  peer-ip  tunnel-i/f
-----
vpn-to-siteB  5    active 100.1.1.1 200.1.1.1 tunnel.41
```

## 2.9.5 Generic Routing Encapsulation

A Generic Routing Encapsulation (GRE) tunnel connects two endpoints (a firewall and another appliance) in a point-to-point, logical link. The firewall can terminate GRE tunnels; you can route or forward packets to a GRE tunnel. GRE tunnels are simple to use and often the tunneling protocol of choice for point-to-point connectivity, especially to the services in the cloud or to partner networks.

Create a GRE tunnel when you want to direct packets that are destined for an IP address to take a certain point-to-point path—for example, to a cloud-based proxy or to a partner network. The packets travel through the GRE tunnel (over a transit network such as the internet) to the cloud service while on their way to the destination address. This enables the cloud service to enforce its services or policies on the packets.

The following figure is an example of a GRE tunnel connecting the firewall across the internet to a cloud service.



When the firewall allows a packet to pass (based on a policy match) and the packet egresses to a GRE tunnel interface, the firewall adds GRE encapsulation; it doesn't generate a session. The firewall does not perform a Security policy rule lookup for the GRE-encapsulated traffic, so you don't need a Security policy rule for the GRE traffic that the firewall encapsulates. However, when the firewall receives GRE traffic, it generates a session and applies all of the policies to the GRE IP header in addition to the encapsulated traffic. The firewall treats the received GRE packet like any other packet. Therefore:

- If the firewall receives the GRE packet on an interface that has the same zone as the tunnel interface associated with the GRE tunnel (such as tunnel.1), the source zone is the same as the destination zone. By default, traffic is allowed within a zone (intrazone traffic), so the ingress GRE traffic will be allowed by default.
- However, if you configured your own intrazone Security policy rule to deny such traffic, you must explicitly allow GRE traffic.
- Likewise, if the zone of the tunnel interface associated with the GRE tunnel (for example, tunnel.1) is a different zone from that of the ingress interface, you must configure a Security policy rule to allow the GRE traffic.

The firewall encapsulates the tunneled packet in a GRE packet, and so the additional 24 bytes of GRE header automatically result in a smaller MSS in the MTU. If you don't change the IPv4 MSS adjustment size for the interface, the firewall reduces the MTU by 64 bytes by default (40 bytes of IP header + 24 bytes of GRE header). This means that if the default MTU is 1,500 bytes, the MSS will be 1,436 bytes ( $1,500 - 40 - 24 = 1,436$ ). For example, if you configure an MSS adjustment size of 300 bytes, the MSS will only be 1,176 bytes ( $1,500 - 300 - 24 = 1,176$ ).

The firewall does not support routing a GRE or IPsec tunnel to a GRE tunnel, but you can route a GRE tunnel to an IPsec tunnel. Additionally:

- A GRE tunnel does not support QoS.
- The firewall does not support a single interface acting as both a GRE tunnel endpoint and a decryption broker.
- GRE tunneling does not support NAT between the GRE tunnel endpoints.

### 2.9.6 One-to-one and one-to-many tunnels

Palo Alto Networks supports the following VPN deployments:

- **Site-to-site VPN:** This deployment provides a simple VPN that connects a central site and a remote site. This is also commonly referred to as a hub-and-spoke VPN that connects a central (gateway) site with multiple remote (branch) sites.
- **Remote-user-to-site VPN:** This deployment provides an endpoint client to use the GlobalProtect agent for a secure remote user access connection through the firewall gateway.
- **Large scale VPN (LSVPN):** This deployment uses the Palo Alto Networks GlobalProtect LSVPN. It provides a scalable mechanism to provide a hub-and-spoke VPN for up to 1,024 branch offices.

### 2.9.7 Determine when to use proxy IDs

#### Symptom

When configuring IPsec VPNs, Proxy IDs are a requirement with a peer that supports policy-based VPNs.

Sometimes, multiple local and remote subnets need to communicate over VPN for the same peer. If peer side is a policy-based VPN, you need to set up multiple proxy IDs on the Palo Alto firewall Tunnel configuration to match with the peer's policies.

Even with the correct configuration, the traffic might fail because of the way in which proxy IDs are stored in the data plane (DP). This article highlights the best practices to be used when configuring multiple Proxy IDs with the same peer, which are for overlapping subnets.

#### Environment

- Any PAN-OS
- Palo Alto Networks Firewall
- IPSEC VPN configured with Proxy IDs

## Cause

When multiple Proxy IDs are configured, the naming of the Policy IDs is important because the order of the proxy ID matching depends on the string order of the proxy ID name.

Example:

Let's say, there are four Proxy IDs configured under the tunnel configuration, as follows:

```
TestProxyID-1      : Local = 10.1.1.0/24,   Remote = 192.168.30.0/24
ProxyID-10_8_0_0  : Local = 10.8.1.0/24,   Remote = 192.168.30.0/24
proxy-id-10_123_0_0 : Local = 10.123.1.0/24, Remote = 192.168.30.0/24
AllNetworks       : Local = 10.0.0.0/8,    Remote = 192.168.30.0/24
```

When the proxy IDs are stored in a DP, they are sorted using String Comparison (ASCII sorting).

Using the above, the string sort order for the above proxy ID names would be as follows:

```
AllNetworks       : Local = 10.0.0.0/8,    Remote = 192.168.30.0/24
proxy-id-10_123_0_0 : Local = 10.123.1.0/24, Remote = 192.168.30.0/24
ProxyID-10_8_0_0  : Local = 10.8.1.0/24,   Remote = 192.168.30.0/24
TestProxyID-1     : Local = 10.1.1.0/24,   Remote = 192.168.30.0/24
```

The IPSEC Security SA's will be stored in this order in the DP. This will affect traffic processing because when a certain traffic needs to be encrypted using one of the proxy IDs, it will look from top to bottom for the first matching proxy ID.

In the above example, even though the **"AllNetworks"** proxy ID is defined on the bottom in the configuration, it will be the first in order in DP.

In the above example, if any traffic is going from source 10.123.1.0/24 via this IPSEC tunnel to a remote IP, it will not be sent via **"proxy-id-10\_123\_0\_0"** but via **"AllNetworks"**. So this might fail on the remote side that is checking incoming traffic against the proxy IDs.

## Resolution

For proxy IDs with overlapping subnets, define the proxy ID names so that a more specific proxy ID name is above the broader Proxy ID name, as per String Sorting.

## 2.9.8 References

- Tunnel Interface, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/site-to-site-vpn-concepts/tunnel-interface>
- Site-to-Site VPN with Static and Dynamic Routing, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/site-to-site-vpn-quick-configs/site-to-site-vpn-with-static-and-dynamic-routing>
- VPN Deployments, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/vpn-deployments>
- Site-to-Site VPN Overview, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/site-to-site-vpn-overview>
- Large Scale VPN (LSVPN), <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/large-scale-vpn-lsvpn>
- Network > Network Profiles > Monitor, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/network/network-network-profiles/network-network-profiles-monitor.html>
- Test VPN Connectivity, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/set-up-site-to-site-vpn/test-vpn-connectivity.html>
- Configure Multiple Proxy IDs in VPN Tunnel with Overlapping Subnet Ranges, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000PLTICAO>

## 2.10 Configure service routes

Configure service routes globally for the firewall. Any virtual system that does not have a service route configured for a particular service inherits the interface and IP address that are set globally for that service.

### 2.10.1 Default

Perform the following tasks to configure static routes or a default route for a virtual router on the firewall.

**Step 1:** Configure a static route.

- Select **Network > Virtual Router** and then select the virtual router you are configuring—for example, **default**.
- Select the **Static Routes** tab.
- Select **IPv4** or **IPv6**, depending on the type of static route you want to configure.
- Add a **Name** for the route. The name must start with an alphanumeric character, and it can contain a combination of alphanumeric characters, underscores (\_), hyphens (-), dots (.), and spaces. Beginning with PAN-OS 10.0.8, the name can be a maximum of 63 characters.
- For **Destination**, enter the route and netmask (for example, 192.168.2.2/24 for an IPv4 address or 2001:db8:123:1::1/64 for an IPv6 address). If you're creating a default route, enter the default route (0.0.0.0/0 for an IPv4 address or ::/0 for an IPv6 address). Alternatively, you can create an address object of type IP Netmask.

- (Optional) For **Interface**, specify the outgoing interface for packets to use to go to the next hop. Use this for stricter control over which interface the firewall uses rather than the interface in the route table for the next hop of this route.
- For **Next Hop**, select one of the following:
  - **IP Address:** Enter the IP address (for example, 192.168.56.1 or 2001:db8:49e:1::1) when you want to route to a specific next hop. You must enable IPv6 on the interface (when you configure Layer 3 interfaces) to use an IPv6 next hop address. If you're creating a default route for **Next Hop**, you must select **IP Address** and enter the IP address for your internet gateway (for example, 192.168.56.1 or 2001:db8:49e:1::1). Alternatively, you can create an address object of type IP Netmask. The address object must have a netmask of /32 for IPv4 or /128 for IPv6.
  - **Next VR:** Select this option and then select a virtual router if you want to route internally to a different virtual router on the firewall.
  - **FQDN:** Enter an FQDN or select an address object that uses an FQDN. You can also create a new address object of type FQDN.

If you use an FQDN as a static route next hop, the FQDN must resolve to an IP address that belongs to the same subnet as the interface you configured for the static route; otherwise, the firewall rejects the resolution and the FQDN remains unresolved.

The firewall uses only one IP address (from each IPv4 or IPv6 family type) from the DNS resolution of the FQDN. If the DNS resolution returns more than one address, the firewall uses the preferred IP address that matches the IP family type (IPv4 or IPv6) configured for the next hop. The preferred IP address is the first address the DNS server returns in its initial response. The firewall retains this address as the preferred one as long as the address appears in subsequent responses, regardless of its order.

**Discard:** Select to drop the packets that are addressed to this destination.

- **None:** Select if there is no next hop for the route. For example, a point-to-point connection does not require a next hop because there is only one way for the packets to go.
- Enter an **Admin Distance** for the route to override the default administrative distance set for the static routes on this virtual router (range is 10 to 240; default is 10).
- Enter a **Metric** for the route (range is 1 to 65,535).

**Step 2:** Choose where to install the route.

Select the RIB into which you want the firewall to install the static route:

- **Unicast:** Install the route in the unicast route table. Choose this option if you want the route to be used only for unicast traffic.
- **Multicast:** Install the route in the multicast route table (available for the IPv4 routes only). Choose this option if you want the route to be used only for multicast traffic.
- **Both:** Install the route in the unicast and multicast route tables (available for the IPv4 routes only). Choose this option if you want either unicast or multicast traffic to use the route.
- **No Install:** Do not install the route in either route table.

**Step 3:** (Optional) If your firewall model supports bidirectional forwarding detection (BFD), you can apply a BFD Profile to the static route so that if the static route fails, the firewall removes the route from the RIB and FIB and uses an alternative route. The default is None.

**Step 4:** Click **OK** twice.

**Step 5:** Commit the configuration.

### 2.10.2 Custom

When a firewall is enabled for multiple virtual systems, the virtual systems inherit the global service and service route settings. For example, the firewall can use a shared email server to originate email alerts to all the virtual systems. In some scenarios, you will want to create different service routes for each virtual system.

One use case for configuring service routes at the virtual system level is an ISP that needs to support multiple individual tenants on a single Palo Alto Networks firewall.

### 2.10.3 Destination

On the **Global** tab, when you click **Service Route Configuration > Customize**, the **Destination** tab appears. Destination service routes are available under the Global tab only (not the **Virtual Systems** tab) so that the service route for an individual virtual system cannot override any route table entries that are not associated with that virtual system. You can use a destination service route to add a customized redirection of a service that is not supported on the customized list of services. A destination service route is a way to set up routing to override the FIB route table. Any settings in the destination service routes override the route table entries. They could be related or unrelated to any service.

The Destination tab is for the following use cases:

- When a service does not have an application service route
- Within a single virtual system, when you want to use multiple virtual routers or a combination of a virtual router and management port

DESTINATION SERVICE ROUTE SETTINGS	DESCRIPTION
Destination	Enter the Destination IP address. An incoming packet with a destination IP address that matches this address will use as its source the Source Address you specify for this service route.
Source Interface	To limit the drop-down options for Source Address, select a Source Interface. Selecting Any causes all IP addresses on all of the interfaces to be available in the Source Address drop-down list. Selecting MGT causes the firewall to use the MGT interface for the service route.
Source Address	Select the Source Address for the service route; this address will be used for packets returning from the destination. You do not need to enter the subnet for the destination address.

## 2.10.4 Custom routes for different virtual systems versus destination routes

### Virtual Systems

When a firewall is enabled for multiple virtual systems, the virtual systems inherit the global service and service route settings. For example, the firewall can use a shared email server to originate email alerts to all virtual systems. In some scenarios, you'd want to create different service routes for each virtual system.

One use case for configuring service routes at the virtual system level is if you are an ISP who needs to support multiple individual tenants on a single Palo Alto Networks firewall. Each tenant requires custom service routes to access services, such as DNS, Kerberos, LDAP, NetFlow, RADIUS, TACACS+, Multi-Factor Authentication, email, SNMP trap, syslog, HTTP, User-ID Agent, VM Monitor, and Panorama (deployment of content and software updates). Another use case is an IT organization that wants to provide full autonomy to the groups that set servers for services. Each group can have a virtual system and define its own service routes.

You can select a virtual router for a service route in a virtual system; you cannot select the egress interface. After you select the virtual router and the firewall sends the packet from the virtual router, the firewall selects the egress interface based on the destination IP address. Therefore, if a virtual system has multiple virtual routers, packets to all of the servers for a service must egress out of only one virtual router. A packet with an interface source address may egress a different interface, but the return traffic would be on the interface that has the source IP address, creating asymmetric traffic.

### Destination Routes

On the **Global** tab, when you click **Service Route Configuration > Customize**, the **Destination** tab appears. Destination service routes are available under the Global tab only (not the **Virtual Systems** tab) so that the service route for an individual virtual system cannot override the route table entries that are not associated with that virtual system. You can use a destination service route to add a customized redirection of a service that is not supported on the customized list of services. A destination service route is a way to set up routing to override the FIB route table. Any settings in the destination service routes override the route table entries. They could be related or unrelated to any service.

The Destination routes use cases are similar to the ones discussed in the "Destination" section.

## 2.10.5 How to verify service routes

Configure service routes globally for the firewall. Any virtual system that does not have a service route configured for a particular service inherits the interface and IP address that are set globally for that service.



## 2.10.6 References

- Configure a Static Route:  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/static-routes/configure-a-static-route>
- Service Routes:  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/service-routes>

## 2.11 Configure application-based QoS

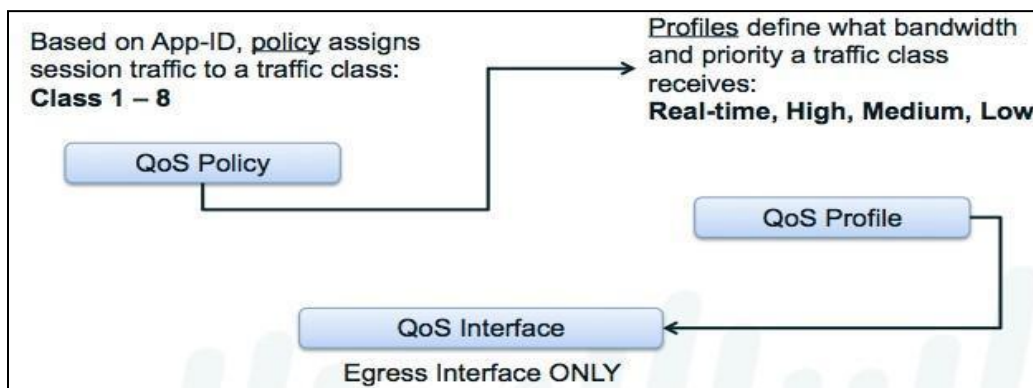
QoS is a set of technologies that works on a network to guarantee its ability to run high-priority applications and traffic dependably with shared network capacity. QoS technologies achieve this by providing differentiated handling and capacity allocation to specific flows in network traffic, which enables the network administrator to assign the order in which traffic is handled and the amount of bandwidth provided to traffic.

### 2.11.1 Enablement requirements

#### QoS

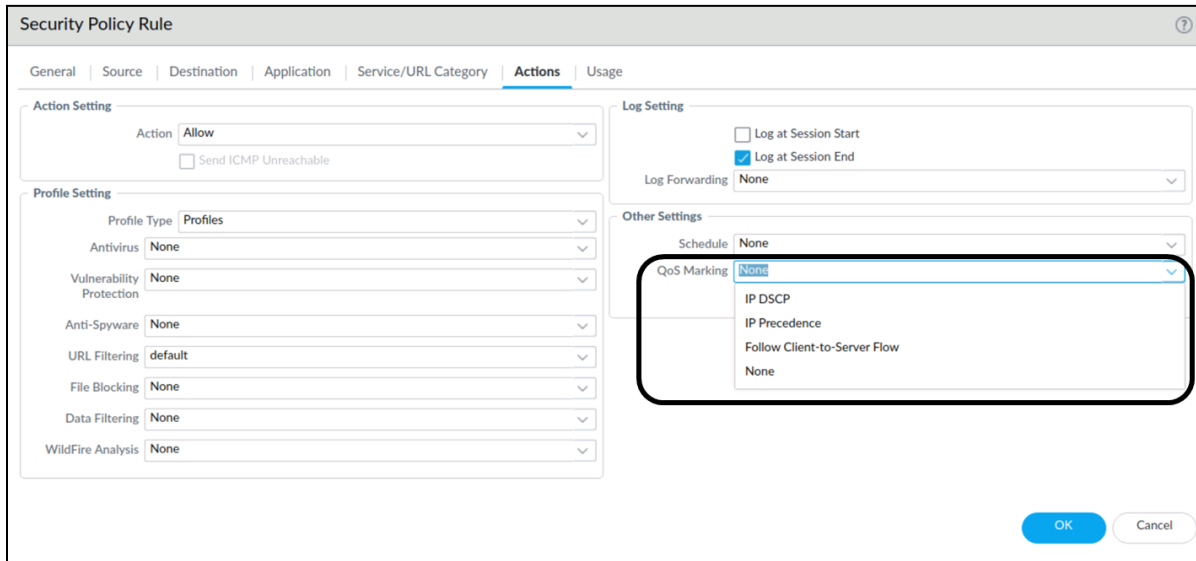
Palo Alto Networks QoS provides an “application-aware” QoS service that can be driven by the traffic’s App-ID. The firewall’s QoS implementation is a self-contained system local to the firewall, which can consider existing QoS packet markings but does not act directly on them. Traffic is evaluated against QoS policy rules, including existing QoS packet markings, App-ID, and other matching conditions, to assign a traffic classification value of 1 through 8. These values are the basis for QoS decision making. QoS traffic control is limited to egress traffic for the configured interface(s) only. Ingress traffic cannot be managed.

The interrelationship between the QoS policies, traffic classes, QoS Profiles, and interfaces is displayed in the following figure:



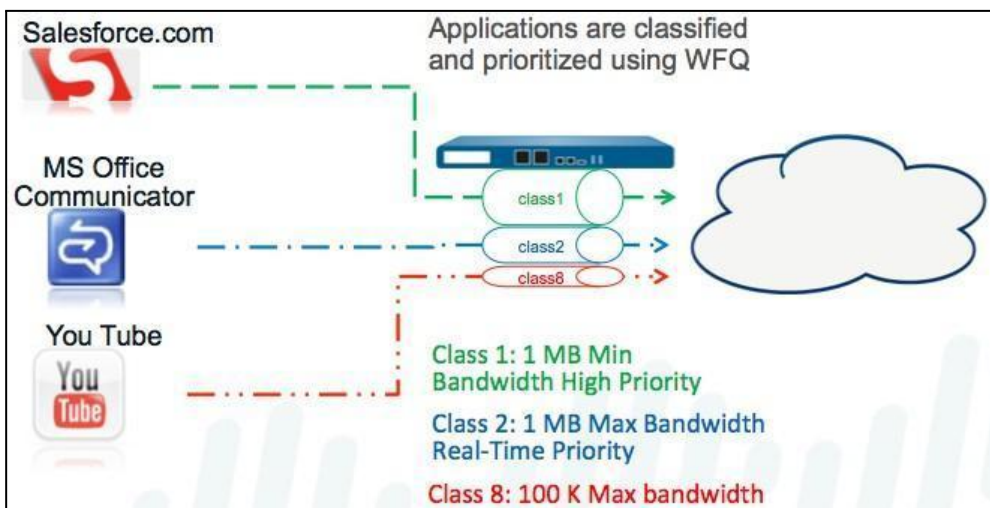
## 2.11.2 QoS policy rule

Use the QoS Marking field when setting up a Security policy rule to write QoS marking into packet headers. This applies to any traffic that the Security policy rule processes. Note that this marking is not directly related to QoS processing in the firewall.



QoS implementation on a Palo Alto Networks firewall begins with three primary configuration components that support a full QoS solution: a QoS policy, a QoS Profile, and a configuration of the QoS egress interface. Each option in the QoS configuration task facilitates a broader process that optimizes and prioritizes the traffic flow and allocates and ensures bandwidth according to configurable parameters.

QoS policies assign classes (1 to 8) to the traffic that matches the policy conditions. The PAN-OS QoS functionality can use App-ID for specific bandwidth reservation.



### 2.11.3 Add a Differentiated Services Code Point/ToS component

A Differentiated Services Code Point (DSCP) is a packet header value that can be used to request (for example) high-priority or best-effort delivery for traffic. Session-based DSCP classification allows you to both honor DSCP values for incoming traffic and mark a session with a DSCP value as session traffic exits the firewall. This enables all inbound and outbound traffic for a session to receive continuous QoS treatment as it flows through the network. For example, inbound return traffic from an external server can now be treated with the same QoS priority that the firewall initially enforced for the outbound flow based on the DSCP value the firewall detected at the beginning of the session. Network devices between the firewall and end user will also then enforce the same priority for the return traffic (and any other outbound or inbound traffic for the session).

Different types of DSCP markings indicate different levels of service:

Completing this step enables the firewall to mark traffic with the same DSCP value that was detected at the beginning of a session (in this example, the firewall would mark return traffic with the DSCP AF11 value). While configuring QoS allows you to shape traffic as it egresses the firewall, enabling this option in a Security rule allows the other network devices intermediate to the firewall and the client to continue to enforce priority for DSCP-marked traffic.

**Expedited Forwarding (EF):** Can be used to request low loss, low latency, and guaranteed bandwidth for traffic. Packets with EF codepoint values are typically guaranteed the highest-priority delivery.

**Assured Forwarding (AF):** Can be used to provide reliable delivery for applications. Packets with AF codepoint indicate a request for the traffic to receive higher priority treatment than what the best-effort service provides (although packets with an EF codepoint continue to take precedence over those with an AF codepoint).

**Class Selector:** Can be used to provide backward compatibility with network devices that use the IP precedence field to mark priority traffic.

**IP Precedence (ToS):** Can be used by legacy network devices to mark priority traffic (the IP precedence header field was used to indicate the priority for a packet before the introduction of the DSCP classification).

**Custom Codepoint:** Can be used to match to traffic by entering a codepoint name and binary value.









For example, select **AF** to ensure that traffic marked with an AF codepoint value has higher priority for reliable delivery over applications marked to receive lower priority. To enable session-based DSCP classification, start by configuring QoS based on DSCP marking detected at the beginning of a session. You can then continue to enable the firewall to mark the return flow for a session with the same DSCP value used to enforce QoS for the initial outbound flow.

## 2.11.4 QoS Profile

QoS Profiles describe the priority to be given to the specified traffic when the interface becomes constrained. As priority decreases, more packets are randomly dropped until the constraint is cleared. The number of packets dropped is determined by their assigned priority. A real-time priority setting means that no packet dropping will be performed. High-, medium-, and low-priority settings indicate that greater levels of random packet dropping will be performed during movement down the scale. No packets will be dropped until the egress traffic on the managed interface becomes constrained, meaning that outbound traffic queues for the interface will fill faster than they can be emptied.

Profiles also specify the maximum bandwidth enforcement that is always applied. Bandwidth that is configured as the maximum limit can be used by all of the traffic until the interface becomes constrained. After an interface is constrained, sessions might receive no more than their guaranteed bandwidth.

QoS Profiles prioritize specified traffic. The following figure shows the four possible priority values:

NAME	GUARANTEED EGRESS	MAXIMUM EGRESS	PRIORITY
 default			
 class1			real-time
 class2			high
 class3			high
 class4			medium
 class5			medium
 class6			low
 class7			low
 class8			low

## 2.11.5 Determine how to control bandwidth use on a per-application basis

Voice and video traffic is particularly sensitive to the measurements that the QoS feature shapes and controls, especially latency and jitter. For voice and video transmissions to be audible and clear, voice and video packets cannot be dropped, delayed, or delivered inconsistently. A best practice for voice and video applications, in addition to guaranteeing bandwidth, is to guarantee priority to voice and video traffic.

## 2.11.6 Use QoS to monitor bandwidth utilization

QoS bandwidth management allows you to control traffic flows on a network so that traffic does not exceed network capacity and result in network congestion. This also allows you to allocate bandwidth for certain types of traffic and for applications and users. With QoS, you can enforce bandwidth for traffic on a narrow or broad scale. A QoS Profile rule allows you to set bandwidth limits for individual QoS classes and the total combined bandwidth for all eight QoS classes. As part of the steps to configure QoS, you can attach the QoS Profile rule to a physical interface to enforce bandwidth settings on the traffic exiting that interface; the individual QoS class settings are enforced on traffic matching QoS policy rules. The overall bandwidth limit for the profile can be applied to all of the cleartext traffic, specific cleartext traffic originating from source interfaces and source subnets, all of the tunneled traffic, and individual tunnel interfaces. You can add multiple profile rules to a single QoS interface to apply varying bandwidth settings to the traffic exiting that interface.

Egress guaranteed and egress max both support QoS bandwidth settings.

### Egress Guaranteed

**Egress guaranteed specifies** the amount of bandwidth guaranteed for matching traffic. When the egress-guaranteed bandwidth is exceeded, the firewall passes traffic on a best-effort basis. Bandwidth that is guaranteed but is unused continues to remain available for all of the traffic. Depending on the QoS configuration, you can guarantee bandwidth for a single QoS class, for all or some cleartext traffic, and for all or some tunneled traffic.

**Example:** Class 1 traffic has 5Gbps of egress guaranteed bandwidth, which means that 5Gbps is available but is not reserved for Class 1 traffic. If Class 1 traffic does not use or only partially uses the guaranteed bandwidth, the remaining bandwidth can be used by other classes of traffic. However, during high-traffic periods, 5Gbps of bandwidth is absolutely available for Class 1 traffic. During these periods, any Class 1 traffic that exceeds 5Gbps is passed on a best-effort basis.

### Egress Max

**Egress max specifies** the overall bandwidth allocation for matching traffic. The firewall drops the traffic that exceeds the egress max limit set. Depending on the QoS configuration, you can set a maximum bandwidth limit for a QoS class, for all or some cleartext traffic, for all or some tunneled traffic, and for all of the traffic exiting the QoS interface.

## 2.11.6 References

- QoS Policy Rule,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/quality-of-service/enforce-qos-based-on-dscp-classification>
- Quality of Service,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/quality-of-service>
- QoS Bandwidth Management,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/quality-of-service/qos-concepts/qos-bandwidth-management>
- Use Case: QoS for Voice and Video Applications,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/quality-of-service/qos-use-cases/use-case-qos-for-voice-and-video-applications>

## Domain 3: Deploy and Configure Features and Subscriptions

### 3.1 Configure App-ID

#### 3.1.1 Create security rules with App-ID

##### Security Policy Overview

The firewall does not allow any traffic to flow from one zone to another unless a Security policy rule allows it. When a packet enters a firewall interface, the firewall matches the attributes in the packet against the Security policy rules to determine whether to block or allow the session based on attributes, such as the source and destination security zone, the source and destination IP address, the application, the user, and the service. The firewall evaluates all incoming traffic against the Security policy rulebase from left to right and from top to bottom. Then, the firewall takes the action specified in the first Security rule that matches (for example, whether to allow, deny, or drop the packet). Processing occurs from top to bottom, so you must order the rules in your Security policy rulebase so that the more specific rules are at the top of the rulebase and the more general rules are at the bottom. This ensures that the firewall enforces policy as expected.

##### Configuring Security Rules

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

Security policy rules are matched from top down. Up to two processing steps are in each Security policy match. Step 1 confirms that a match has been made based on the matching conditions in the Security policy. If a match is found in Step 1, the traffic is logged (based on that policy rule's configuration) and the chosen action (deny, allow, drop, reset) is performed. Once processing is complete, no further matching takes place in the Security policy rulebase.

##### Security Policy: Allow

If the action is Allow, Step 2 of the policy is evaluated. Step 2 is the application of the configured Security Profiles. In Step 2, the content of the sessions is scanned for various threat signatures. In this step, URLs can also be scanned for unauthorized destinations and files can be scanned for malware.

If Panorama device groups are used to push the Security policy to one or more firewalls, the Security policy list is expanded to include rules before (pre) and after (post) the local firewall rules. Panorama rules are merged with the local firewall policies in the position chosen during Panorama rule creation. Panorama-supplied rules are read-only to the local firewall administrators. The Security policy rule list displayed in the following screenshot is for a local administrator logged directly into a managed firewall.

	NAME	Source	Destination	APPLICATION	SERVICE	ACTION	PROFILE	OPTIONS
		ZONE	ZONE					
Pre rules from Panorama	1 Block_Bad_IPs_Inbound	Internet	Extranet Users_Net	any	application-default	Deny	none	
	2 Block_Bad_IPs_Outbound	Extranet Internet	Internet	any	application-default	Deny	none	
Local rules created directly in the firewall	3 Local-Allow Facebook	Internet	Internet	facebook mqtt rtcp rtp-base ssl stun web-browsing	application-default	Allow		
	4 Users_to_Extranet	Users_Net	Extranet	any	any	Allow		
	5 Extranet_to_Internet	Extranet	Internet	any	application-default	Allow		
	6 Extranet_to_Users_Net	Extranet	Users_Net	any	application-default	Allow	none	
	7 Danger_Traffic	Danger	any	any	application-default	Allow		
	8 Allow-Internet-Access	Users_Net	Internet	any	application-default	Allow		
Default rules from Panorama	9 intrazone-default	any	(intrazone)	any	any	Allow		
	10 interzone-default	any	any	any	any	Deny	none	

Security policy should use App-ID for match criteria rather than only services (ports).

At the end of the list are two default policy rules, one for an intrazone Allow and one for an interzone Deny. Together, they implement the default security behavior of the firewall to block interzone traffic and allow intrazone traffic. By default, traffic logging is disabled in both rules.

Security policy rules in PAN-OS are configured by type: universal (default), interzone, and intrazone. (All policy rules – regardless of type – are evaluated top down, first match, then exit.) The universal type covers both interzone and intrazone.

The screenshot shows the 'Security Policy Rule' configuration interface. The 'Rule Type' dropdown menu is open, displaying three options: 'universal (default)', 'intrazone', and 'interzone'. The 'universal (default)' option is selected and highlighted. The interface includes tabs for 'General', 'Source', 'Destination', 'Application', 'Service/URL Category', and 'Actions'. Other visible fields include 'Name', 'Description', 'Tags', 'Group Rules By Tag' (set to 'None'), and 'Audit Comment'. A link for 'Audit Comment Archive' is located at the bottom of the form.

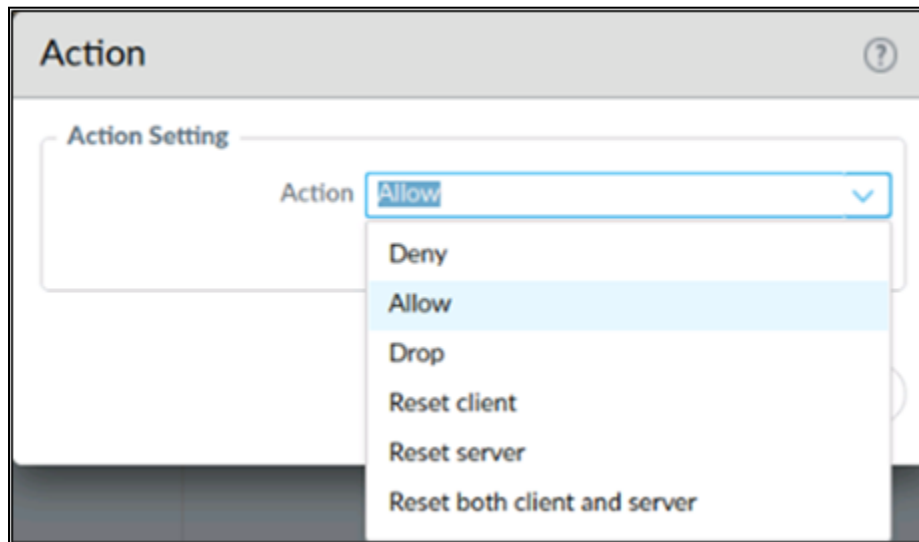
The Security policy **Rule Type** selects the type of traffic to which the policy applies.

Throughput performance does not change based on how quickly a match is made. Evaluation is top down, first match, then exit; therefore, exceptions to policy rules must appear before the general policy. Beyond this policy, order is based on administrative preference. Use tags, a policy search bar, or a global find to navigate quickly to the policy or policy rules needed for troubleshooting and for moves, additions, changes, deletions, and clones.



## Security Policy: Deny

The Deny action is a legacy setting from the prior versions of PAN-OS when denying traffic was the only way to stop traffic. Before PAN-OS 7, the software referenced the App-ID database to find the preferred method of stopping traffic for the matching session's application, which ranged from blocking to reset. These choices have now been added directly to the list of action choices available. Firewall administrators can now choose the desired blocking action directly, or they can continue to rely on the Palo Alto Networks specification by choosing **Deny**.



## App-ID vs. Port-Based Security

Security policy rules that evaluate based on protocol type and port numbers are not accurate enough to control application access effectively through the firewall. Many applications use alternative or even multiple port numbers, making their detection even more difficult. For example, allowing TCP port 80 provides access for all the web-based applications and their associated vulnerabilities.

The Palo Alto Networks App-ID enables positive application identification, regardless of port usage. App-ID allows you to enable safe access to only those applications that you want users to reach. This practice reduces the attack surface by eliminating unauthorized applications and their potential vulnerabilities.

## Note About Using App-ID

Applications often use non-standard ports for communication; therefore, a traffic enforcement technology based only on port numbers does not provide security administrators enough control over the actual application traffic entering their organizations. App-ID identifies applications based primarily on packet contents and not on port numbers and provides a much higher level of capability. When you use Palo Alto Networks firewalls, your Security policy rules should use App-ID, not port numbers, as the selection criteria.

### 3.1.2 Convert port and protocol rules to App-ID rules

#### Moving from Port-Based to App-ID Security

Moving from a port-based Security policy to an application-based Security policy might seem like a daunting task. However, the security risks of staying with a port-based policy far outweigh the effort required to implement an application-based policy. And, although legacy port-based Security policies may have thousands of rules (many of which have an unknown purpose), a best-practice policy has a streamlined set of rules that aligns with business goals, thus simplifying administration and reducing the chance of error. The fact that the rules in an application-based policy align with your business goals and acceptable use policies lets you quickly scan the policy to understand the reason for every rule.

As with any technology, organizations usually take a gradual approach to implementation with carefully planned deployment phases to make the transition as smooth as possible, with minimum impact to end users. The general workflow for implementing a best practice internet gateway Security policy is as follows:

- **Assess your business and identify what you need to protect:** The first step in deploying a security architecture is to assess your business and identify your most valuable assets—and the greatest threats to those assets. For example, if you are a technology company, your intellectual property is your most valuable asset. In this case, one of your biggest threats would be source code theft.
- **Segment your network by using interfaces and zones:** Traffic cannot flow between zones unless there is a Security policy rule to allow it. One of the easiest defenses against lateral attacker movement is to define granular zones and allow access only to the specific user groups that need to access an application or resource in each zone. By segmenting your network into granular zones, you can prevent an attacker from establishing a communication channel within your network (either via malware or by exploiting legitimate applications), thereby reducing the likelihood of a successful attack on your network.
- **Identify allow list applications:** Before you create an internet gateway best practice Security policy, you must have an inventory of the applications that you want to allow on your network. You must also distinguish between those applications that you administer and officially sanction and those that you want users to be able to use safely. After you identify the applications (including general types of applications) that you want to allow, you can map them to specific best practice rules.
- **Create user groups for access to allow list applications:** After you identify the applications that you plan to allow, you must identify the user groups that require access to each one. Compromising an end user's system is one of the cheapest and easiest ways for an attacker to gain access to your network; so you can greatly reduce your attack surface by allowing access to the applications only to the user groups that have a legitimate business need.

- **Decrypt traffic for full visibility and threat inspection:** You cannot inspect traffic for threats if you cannot see it in cleartext. Today, the SSL/TLS traffic flows account for 40% or more of the total traffic on a typical network, which is precisely why encrypted traffic is a common way for attackers to deliver threats. For example, an attacker may use a web application, such as Gmail, which uses SSL encryption, to email an exploit or malware to employees accessing that application on the corporate network. Or, an attacker may compromise a website that uses SSL encryption to download an exploit or malware silently to site visitors. If you are not decrypting traffic for visibility and threat inspection, you are leaving a large surface open for attack.
- **Create best-practice Security Profiles for the internet gateway:** C2 traffic, common vulnerabilities and exposures (CVEs), drive-by downloads of malicious content, phishing attacks, advanced persistent threats—all are delivered via legitimate applications. To protect against known and unknown threats, you must attach stringent Security Profiles to all of the Security policy Allow rules.
- **Define the initial internet gateway Security policy:** Using the application and user group inventory that you conducted, you can define an initial policy that allows access to all the applications you want to allow by user or user group. The initial policy rulebase that you create must also include rules for blocking known malicious IP addresses, temporary rules to prevent other applications you might not have known about from breaking, and the identification of policy gaps and security holes in your design.
- **Monitor and fine-tune the policy rulebase:** After having the temporary rules in place, you can begin monitoring traffic that matches with them so that you can fine-tune your policy. Temporary rules are designed to uncover unexpected traffic on the network (such as traffic running on non-default ports or traffic from unknown users), so you must assess the traffic matching these rules and adjust your application allow rules accordingly.
- **Remove the temporary rules:** After a monitoring period of several months, you should see less and less traffic hitting the temporary rules. When you reach the point where traffic no longer hits the temporary rules, you can remove them.
- **Maintain the rulebase:** Because applications are dynamic, you must continually monitor your application allow list and adapt the rules to accommodate new applications and to determine how new or modified App-IDs impact your policy. The rules in a best-practice rulebase align with your business goals and leverage policy objects; therefore, adding support for a new sanctioned application or a new or modified App-ID is often as simple as adding or removing an application from an application group or modifying an application filter.

Palo Alto Networks has developed an innovative approach toward securing networks, which identifies all the traffic by applications. This approach replaces the conventional approaches that attempt to control traffic based on port numbers.

## Port-Based Rules

When you transition from a legacy firewall to a Palo Alto Networks NGFW, you inherit many port-based rules that allow any application on the allowed ports. This increases the attack surface because any application can use an open port. Policy Optimizer identifies all the applications seen on any legacy port-based Security policy rule and provides an easy workflow for selecting the applications you want to allow on that rule. You can migrate port-based rules to the application-based allow list rules to reduce the attack surface and safely enable applications on your network. Use Policy Optimizer to maintain the rulebase as you add new applications.

### 3.1.3 Identify the impact of application override to overall firewall functionality

#### Application Override Configuration

To change how the firewall classifies network traffic into applications, you can specify Application Override policy rules. These policy rules attach the specified App-ID to matching traffic and bypass the normal App-ID processing steps in the firewall. This assigned application functions identically like an App-ID supplied application name and can be used in the same way. For example, to control a custom application, you can use an Application Override policy to identify traffic for that application according to zone, source, destination address, port, and protocol. After an Application Override rule has assigned a custom application name to the network traffic, that traffic can be controlled by the firewall through the use of the custom application name in a Security policy rule.

Note that the App-ID bypass characteristic of Application Override also skips essential Content-ID processing, which could result in undetected threats. This feature should only be used for trusted traffic.

### 3.1.4 Create custom apps and threats

To make sure that your internal custom apps do not appear as “unknown traffic,” you need to create a custom app. Creating a custom app allows you to minimize the range of incoming unidentified traffic on your network.

To create a custom app, you must define the following app attributes:

- Characteristics
- Category and subcategory
- Risk
- Timeout
- Port

You also must define the patterns or values that the PAN-OS firewall can use to match to the traffic flows themselves (app signatures).

### 3.1.5 Review App-ID dependencies

You now have simplified workflows to find and manage any application dependencies. These workflows allow you to see application dependencies when you create a new Security policy rule and when performing commits. When a policy does not include all of the application dependencies, you can directly access the associated Security policy rule to add the required applications.

Using these workflows along with Policy Optimizer, you can now easily identify, organize, and resolve application dependencies. You can take advantage of the new workflows by upgrading the Panorama management server to 9.1 and pushing rules to the firewalls.

### 3.1.6 References

- Custom Application and Threat Signatures, <https://docs.paloaltonetworks.com/pan-os/u-v/custom-app-id-and-threat-signatures/custom-application-and-threat-signatures>
- Policies > Application Override, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/policies/policies-application-override>
- Defining Applications, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/objects/objects-applications/defining-applications>
- Policies > Security, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/policies/policies-security>
- Set up a Basic Security Policy, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/getting-started/set-up-a-basic-security-policy>
- Internet Gateway Best Practice Security Policy, <https://docs.paloaltonetworks.com/best-practices/10-2/internet-gateway-best-practices>
- Create Best Practice Security Profiles for the Internet Gateway, <https://docs.paloaltonetworks.com/best-practices/10-2/internet-gateway-best-practices/best-practice-internet-gateway-security-policy/create-best-practice-security-profiles>
- Example Web-Based App-ID Listing, <https://applipedia.paloaltonetworks.com/>
- Create a Custom Application Signature, <https://docs.paloaltonetworks.com/pan-os/u-v/custom-app-id-and-threat-signatures/custom-application-and-threat-signatures/create-a-custom-application-signature>
- Resolve Application Dependencies, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/app-id/use-application-objects-in-policy/resolve-application-dependencies>

## 3.2 Configure GlobalProtect

To implement GlobalProtect, configure the following:

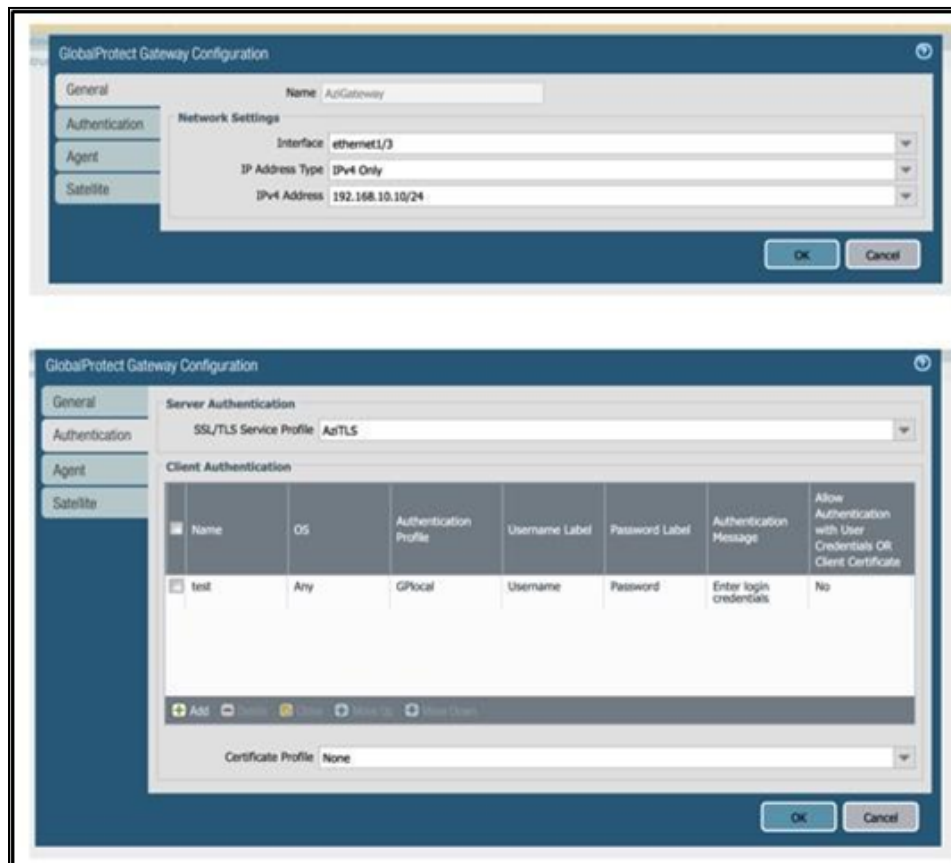
- Download the GlobalProtect client and activate it on the PAN-OS firewall.
- Configure the portal.
- Configure the gateway configuration.
- Set up routing between the trust zones and GlobalProtect clients.
- Configure the Security and NAT policies that permit traffic between the GlobalProtect clients and trust zones.
- Use the GlobalProtect app to connect mobile devices.

### 3.2.1 GlobalProtect licensing

By default, you can deploy the GlobalProtect portals and gateways (without HIP checks) without a license. If you want to use the advanced GlobalProtect features (HIP checks and related content updates, the GlobalProtect mobile app, IPv6 connections, or a GlobalProtect clientless VPN), you will need a GlobalProtect license (subscription) for each gateway.

### 3.2.2 Configure the gateway and the portal

For initial testing, we recommend configuring basic authentication. To authenticate devices with a third-party VPN app, check **Enable X-Auth Support** in the gateway's client configuration. Include the **Group Name** and password for this setting.



## Software Support for GlobalProtect Mobile App 5.0 and PAN-OS 9.0 and later

You can configure a label to identify the physical location of the GlobalProtect gateways and portals by using the CLI or the XML API. The GlobalProtect app displays the location label for the gateway to which users connect. For a clientless VPN, the portal landing page displays the physical location of the portal to which clientless VPN users are logged in.

When end users experience unusual behavior, such as poor network performance, they can provide location information to the support or Help Desk professionals to assist with troubleshooting. They can also use this location information to determine their proximity to the portal or gateway. Based on their proximity, they can evaluate whether they need to switch to a closer portal or gateway.

Refer to the GlobalProtect App 5.0 New Features Guide for more information on the gateway and portal location visibility for end users.

### CLI

Use the following CLI command to specify the physical location of the firewall on which you configured the portal and/or gateway:

```
username@hostname> set deviceconfig setting global-protect location <location>
```

### XML API

Use the following XML API to specify the physical location of the firewall on which you configured the portal and/or gateway:

```
curl -k -F file=@filename.txt -g  
'https://<firewall>/api/?key=<apikey>type=config&action=set&xpath=/config/devices/e  
ntry[@name='<device-name>']/deviceconfig/setting/global-protect&element=<location><  
location-string></location>'
```

#### Where:

**Devices:** Name of the firewall on which you configured the portal and/or gateway

**Location:** Location of the firewall on which you configured the portal and/or gateway

### 3.2.3 GlobalProtect agent

The GlobalProtect agent is a program that runs on the endpoints to protect you by using the same security policies that protect the sensitive resources on your corporate network. (An endpoint can be a desktop computer, laptop, notebook, or smartphone.) You can use the GlobalProtect agent to connect to your corporate network and access your company's internal resources from anywhere in the world.

To install the GlobalProtect agent, download the GlobalProtect VPN client. Choose from the Windows, MacOS, Linux, iOS, and Android client options. The iOS client is available for download from iTunes, and the Android client is available from Google Play.

- To install the agent, install the GlobalProtect Setup Wizard.

Authenticate on the campus VPN network by using DUO two-factor authentication.

### 3.2.4 Differentiate between logon methods

Supported GlobalProtect authentication methods include the following:

- Local authentication
- External authentication
- Client certificate authentication
- Two-factor authentication
- MFA for non-browser-based applications
- SSO

### 3.2.5 Configure clientless VPN

Install a GlobalProtect subscription on the firewall that hosts the clientless VPN from the GlobalProtect portal.

### 3.2.6 HIP

One of the jobs of the GlobalProtect app is to collect information about the host it is running on. The app then submits this host information to the GlobalProtect gateway upon successful connection. The gateway matches this raw host information submitted by the app against any HIP objects and profiles that have been defined. If it finds a match, it generates an entry in the HIP Match log. Additionally, if it finds a HIP Profile match in a policy rule, it enforces the corresponding Security policy.

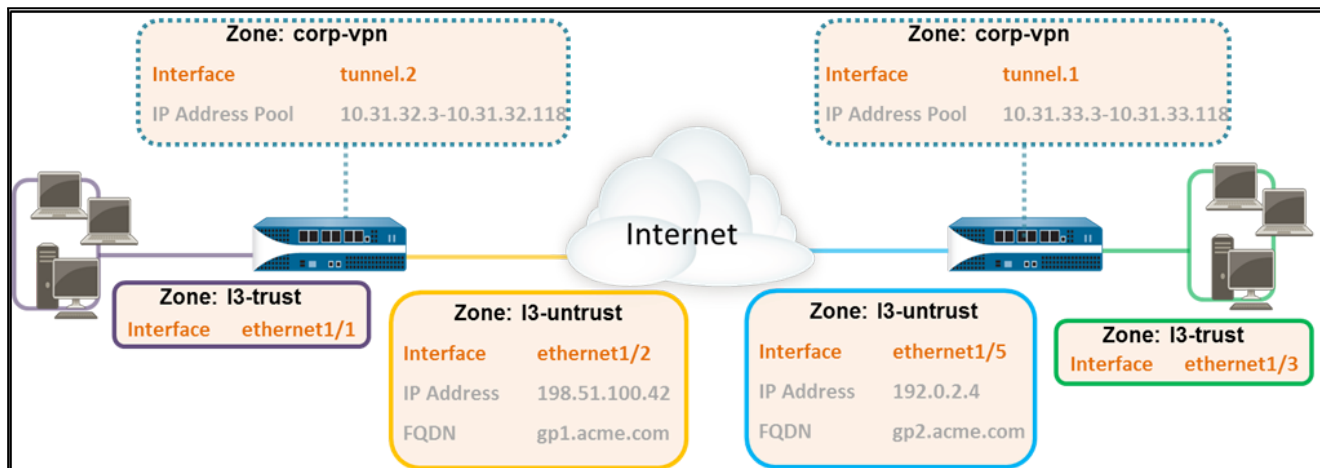
HIP checks are performed when the app connects to the gateway. Subsequent checks are performed hourly while the GlobalProtect agent is connected. The GlobalProtect agent can request an updated HIP report if the previous HIP check has changed. Only the latest HIP report is retained on the gateway per endpoint.

Using HIPs for policy enforcement enables granular security that ensures that the remote hosts accessing critical resources are adequately maintained and adhere with security standards before they are allowed access to network resources. For example, before allowing access to the most sensitive data systems, you might want to ensure that the hosts accessing the data have encryption enabled on their hard drives. You can enforce this policy by creating a Security rule that allows access to the application only if the endpoint system has encryption enabled. In addition, for the endpoints that are not in compliance with this rule, you could create a notification message that alerts users as to why they have been denied access and links them to the file share from where they can access the installation program for the missing encryption software. (Of course, to allow the user to access that file share, you would have to create a corresponding Security rule allowing access to the particular share for hosts with that specific HIP Profile match.)



### 3.2.7 Configure multiple gateway agent profiles

In the GlobalProtect multiple gateway topology below, a second external gateway is added to the configuration. In this topology, you must configure an additional firewall to host the second GlobalProtect gateway. When you add the client configurations to be deployed by the portal, you can also specify different gateways for different client configurations or allow access to all gateways.



If a client configuration contains more than one gateway, the app attempts to connect to all the gateways listed in its client configuration. The app uses priority and response time to determine the gateway to which it will connect. The app only connects to a lower-priority gateway if the response time for the higher-priority gateway is greater than the average response time across all of the gateways.


### 3.2.8 Split tunneling

You can configure split tunnel traffic based on the access route, destination domain, application, and HTTP/HTTPS video streaming application.

The split tunnel capability allows you to conserve bandwidth and route traffic to:

- Tunnel enterprise SaaS and public cloud applications for comprehensive SaaS application visibility and control. This also helps you avoid risks associated with shadow IT in environments where it is not feasible to tunnel all the traffic.
- Send latency-sensitive traffic, such as VoIP, outside the VPN tunnel while sending all other traffic through the VPN for inspection and policy enforcement.
- Exclude the HTTP/HTTPS video-streaming traffic from the VPN tunnel. Video-streaming applications, such as YouTube and Netflix, consume large amounts of bandwidth. By excluding lower-risk video-streaming traffic from the VPN tunnel, you can decrease bandwidth consumption on the gateway.

The split tunnel rules are applied for the Windows and macOS endpoints in the following order:



### 3.2.9 References

- GlobalProtect Overview, <https://docs.paloaltonetworks.com/globalprotect/10-1/globalprotect-admin/globalprotect-overview>
- About Host Information, <https://docs.paloaltonetworks.com/globalprotect/10-1/globalprotect-admin/host-information/about-host-information>
- How to Configure GlobalProtect, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIfbCAK>
- Configure NAT, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/nat/configure-nat>
- Getting Started: Network Address Translation (NAT), <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIIzCAC>
- Global protect multiple gateway configuration, <https://docs.paloaltonetworks.com/globalprotect/10-1/globalprotect-admin/globalprotect-quick-configs/globalprotect-multiple-gateway-configuration>
- Host Information, <https://docs.paloaltonetworks.com/globalprotect/10-1/globalprotect-admin/host-information>
- Split Tunnel Traffic on GlobalProtect Gateways, <https://docs.paloaltonetworks.com/globalprotect/10-1/globalprotect-admin/globalprotect-gateways/split-tunnel-traffic-on-globalprotect-gateways>

### 3.3 Configure decryption

You can configure the firewall to decrypt traffic for visibility, control, and granular security. Decryption policy rules can apply to SSL, including SSL encapsulated protocols (such as IMAP[S], POP3[S], SMTP[S], FTP[S]), and to SSH traffic. SSH decryption can be used to decrypt outbound and inbound SSH traffic to ensure that secure protocols are not being used to tunnel disallowed applications and content.

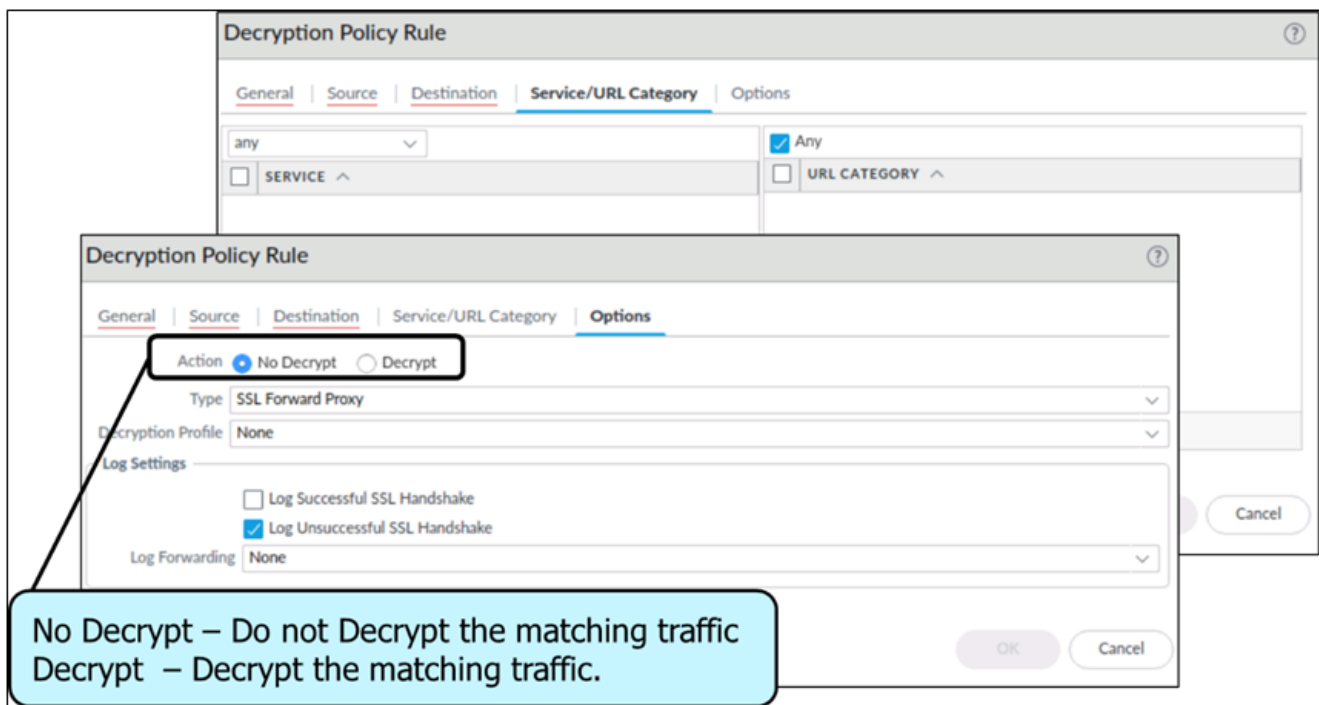
#### 3.3.1 Inbound decryption

##### Configuring Decryption

A Palo Alto Networks firewall can also act as a decryption broker for other external security services. This feature decrypts traffic and forwards it out of the selected interface to a specific security device or service (or chain of devices) that examines the cleartext traffic. The last service in the chain returns the packet to the firewall, which then encrypts it and forwards it to the original destination.

##### Decryption Policies

Ingress traffic decryption is controlled by the decryption policies. Palo Alto Networks firewalls automatically detect encrypted traffic and react by evaluating the decryption policy rules. If a matching policy rule is found, the firewall attempts to decrypt the traffic according to the policy rule's specified decryption action. Normal packet processing resumes afterward.



#### 3.3.2 SSL forward proxy

To configure SSL Forward Proxy decryption, you must set up the certificates required to establish the firewall as a trusted third-party (proxy) to the session between the client and the server.

### 3.3.3 SSL decryption exclusions

The two types of decryption exclusions are predefined exclusions and custom exclusions.

- Predefined decryption exclusions allow applications and services that might break when the firewall decrypts them to remain encrypted. Palo Alto Networks defines the predefined decryption exclusions and delivers updates and additions to the predefined exclusions list at regular intervals as part of the Applications and Threats content update. Predefined exclusions are enabled by default, but you can choose to disable the exclusion as needed.
- You also can create custom decryption exclusions to exclude server traffic from decryption. All the traffic either originating from or destined to the targeted server remains encrypted.

### 3.3.4 SSH proxy

To configure SSH Proxy, you do not need certificates or the key used to decrypt SSH sessions. With SSH decryption enabled, the firewall decrypts SSH traffic and blocks or restricts the SSH traffic based on the decryption policy and Decryption Profile settings. The traffic is re-encrypted as it exits the firewall. Decryption can be performed only on the virtual wire, Layer 2, or Layer 3 interfaces.

### 3.3.5 References

- Decryption Concepts, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/decryption-concept>
- Decryption Best Practices, <https://docs.paloaltonetworks.com/best-practices/10-2/decryption-best-practices>
- IPv6 Support by Feature, <https://docs.paloaltonetworks.com/compatibility-matrix/ipv6-support-by-feature/ipv6-support-by-feature-table>
- Configure SSL Forward Proxy, <https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-admin/decryption/configure-ssl-forward-proxy>

## 3.4 Configure User-ID

### 3.4.1 User-ID agent and agentless

#### User-ID Agent

To map usernames to IP addresses, User-ID agents monitor various sources, such as directory servers. The agents send the user mappings to firewalls, Log Collectors, or Panorama. Each of these appliances can then serve as redistribution points that forward the mappings to other firewalls, Log Collectors, or Panorama. For a firewall (device user identification User-ID agents) or Panorama (Panorama user identification) to collect user mappings, you must configure its connections to the User-ID agents or redistribution points.

#### User-ID Agentless

You can use an agentless User-ID if you have a small-to-medium deployment with 10 or fewer domain controllers or exchange servers and you wish to share the PAN-OS-sourced mappings from AD, Captive Portal, or GlobalProtect with other Palo Alto devices (max 255 devices).

### 3.4.2 User-ID group mapping

The following are the best practices for group mapping in an AD environment:

- If you have a single domain, you only need one group mapping configuration with an LDAP server profile that connects the firewall to the domain controller with the best connectivity. You can add up to four domain controllers to the LDAP server profile for redundancy. Note that you cannot increase redundancy beyond four domain controllers for a single domain by adding multiple group mapping configurations for that domain.
- If you have multiple domains or forests, you must create a group-mapping configuration with an LDAP server profile that connects the firewall to a domain server in each domain or forest. Take steps to ensure unique usernames in separate forests.
  - If you have universal groups, you can create an LDAP server profile to connect to the root domain of the global catalog server on port 3268 or 3269 for SSL, then create another LDAP server profile to connect to the root domain controllers on port 389. This helps ensure that users and group information are available for all the domains and subdomains.
  - Before using group mapping, configure a primary username for user-based Security policies because this attribute identifies users in the policy configuration, logs, and reports.

### 3.4.3 Shared User-ID mapping across virtual systems

You can enable a firewall or virtual system to serve as a data distribution agent that redistributes user mapping information and timestamps associated with authentication challenges. Simply configure the Data Redistribution settings to create an agent that communicates with any firewalls or other devices to share local information.

### 3.4.4 Data redistribution

Every firewall that enforces a user-based policy requires user-mapping information. In a large-scale network, instead of configuring all the firewalls to query the mapping information sources directly, you can streamline resource usage by configuring some firewalls to collect mapping information through redistribution. Redistribution also enables the firewalls to enforce user-based policies when users rely on local sources for authentication (such as regional directory services) but need access to remote services and applications (such as global data center applications). The Data Redistribution feature allows a firewall to be a source of IP user mappings, among other types of data, for any device that is configured to communicate with the agent service of that source firewall or via Panorama.

If you configure an Authentication policy, your firewalls must also redistribute the authentication timestamps that are generated when users authenticate to access applications and services. Firewalls use the timestamps to evaluate the timeouts for the Authentication policy rules. The timeouts allow a user who successfully authenticates to later request services and applications without authenticating again within the timeout periods. The redistribution of timestamps enables you to enforce consistent timeouts across all the firewalls in the network.

Firewalls share user mappings and authentication timestamps as part of the same redistribution flow; you do not have to configure redistribution for each information type separately.

### 3.4.5 User-ID methods

To enforce user- and group-based policies, the firewall must be able to map the IP addresses, in the packets it receives, to usernames. These mappings allow security rules to be enforced appropriately. The following are the different methods of user mapping:

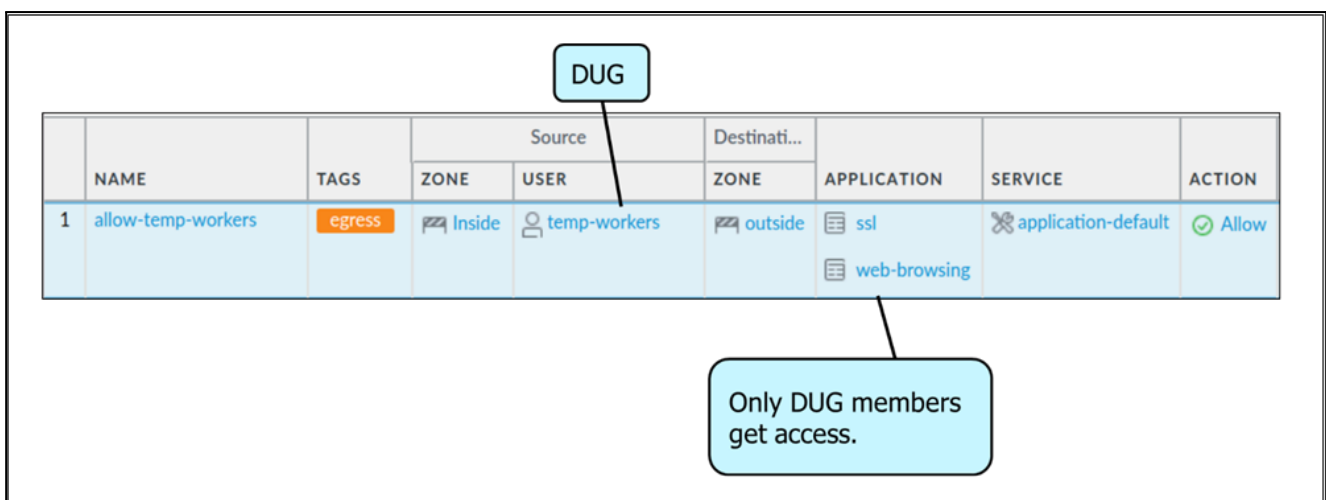
- Server Monitoring
- Port Mapping
- Syslog
- XFF Headers
- Username Header Insertion
- Authentication Policy and Authentication Portal
- GlobalProtect
- XML API
- Client Probing

### 3.4.6 Benefits of using dynamic user groups (DUGs) in policy rules

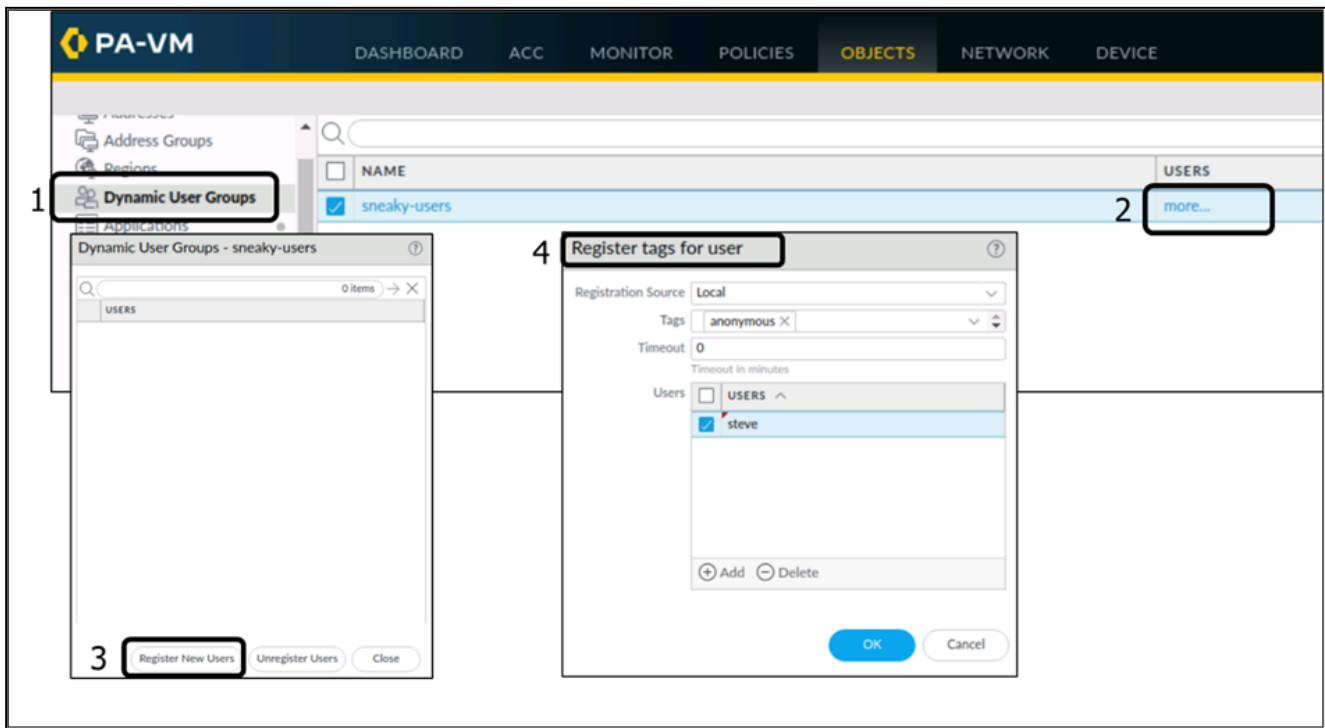
#### Dynamic User Groups

Dynamic user groups (DUGs) control access to the resources that are managed by firewall policies, including the Security policy, the authentication policy, and the decryption policy. DUGs enable you to create policy rules that provide auto-remediation for anomalous user behavior and malicious activity while maintaining user visibility. When you create a policy rule, you can add a DUG to the Source User field as a match criterion. In the past PAN-OS releases, you would have been able to add only a username or a static group name to the Source User field.

You must commit firewall configuration after creating a DUG and adding it to a policy rule. However, you do not have to perform a commit when users are added to or removed from a DUG. User membership in a DUG is dynamic and controlled through the tagging and untagging of usernames. All updates to DUG membership are automatic, and so the use of DUGs instead of a static group (such as an LDAP group) enables you to respond to changes in user behavior or potential threats without needing to make any manual policy changes.



Several methods are available to tag or untag usernames. As shown in the following screenshot, you can manually tag and untag usernames by using the web interface:



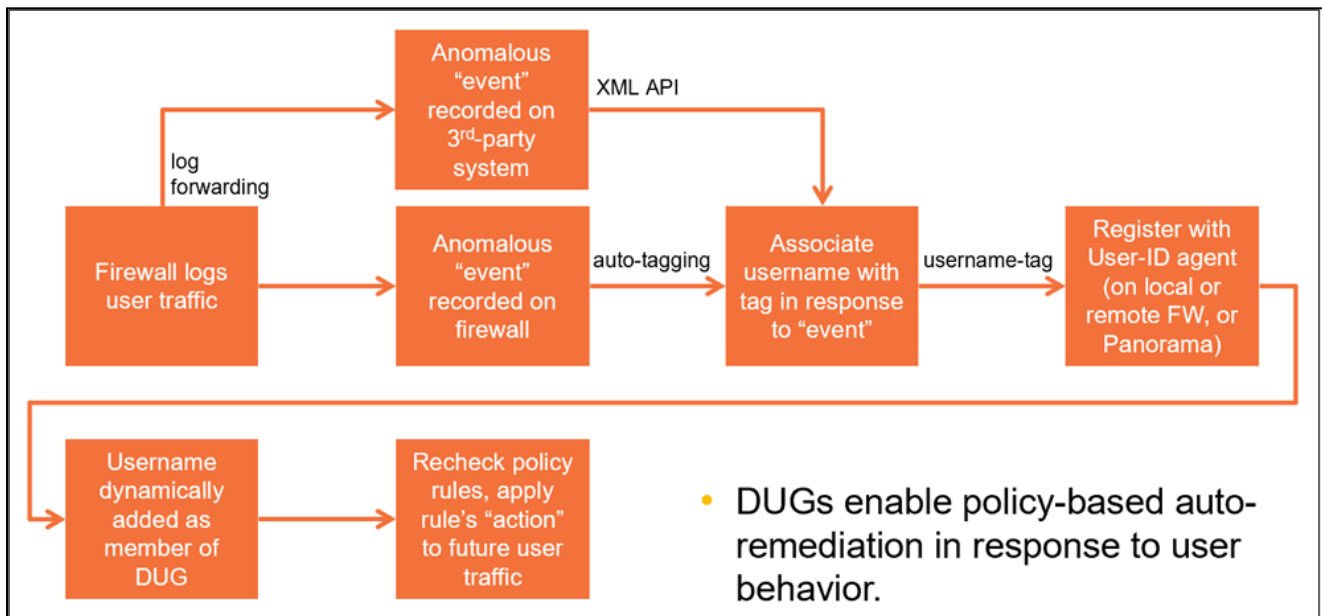
Username tags can also be managed by using the auto-tagging feature in a Log Forwarding Profile. You can also program another utility to invoke the PAN-OS XML API commands to tag or untag usernames. In the web interface, you can use the logical AND or OR operators with the tags to better filter or match against. You can configure a timeout value that determines when a username will be untagged automatically.

### DUG Operation

DUGs enable you to create a Security policy that provides auto-remediation in response to user behavior and activity. Auto-remediation reduces administrative burden by automating the firewall's response to user activity. Auto-remediation by using DUGs also reduces the firewall's response time to malicious activity, which provides better security for the environment.

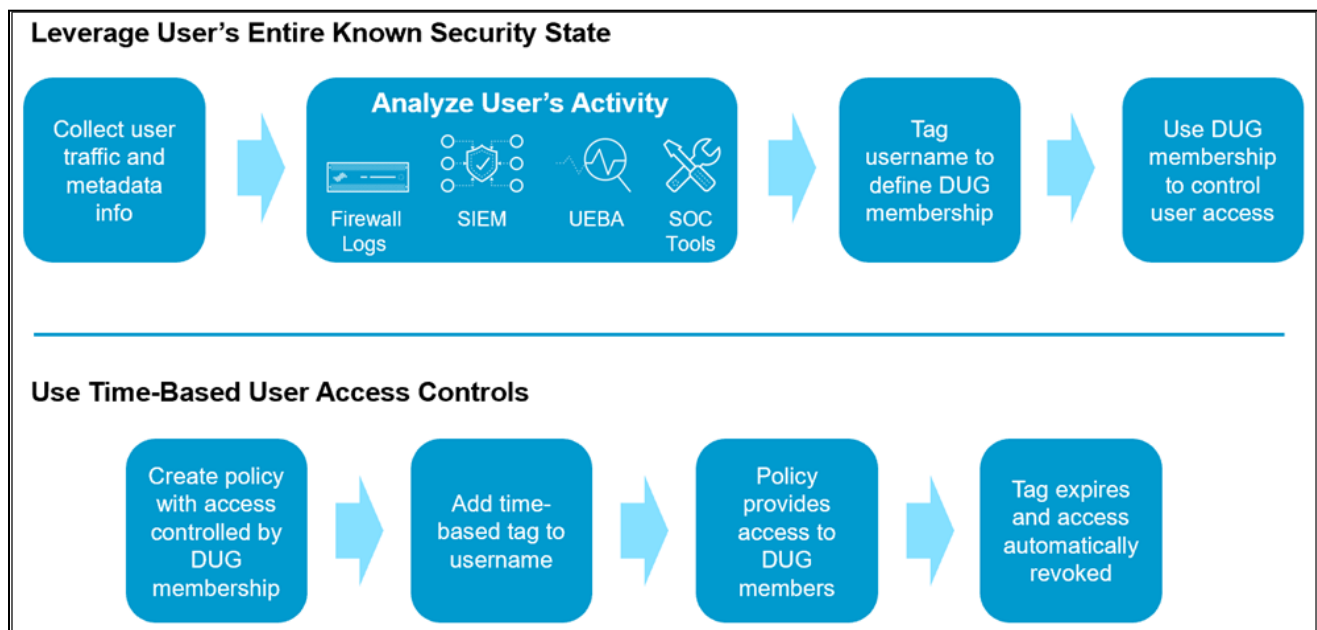
In the following figure, a user's traffic is recorded in the firewall logs. You can analyze these logs directly on the firewall, or you can configure log forwarding to forward the logs to a third-party system for analysis. If logs are being analyzed locally on the firewall, the log-forwarding configuration can invoke a new built-in action that associates a tag with a username based on one or more events in a log. A third-party system can also associate a tag with a username by using the PAN-OS XML API. Username-tag registrations are recorded in and maintained by a User-ID agent.

The firewall uses these username-tag pairs to determine which users are currently members of a DUG. When you configure a DUG, you associate it with one or more tags. Any user who is also associated with a tag configured in a DUG becomes a member of the DUG. The DUG membership is then used to determine the future policy rule matches. For example, a Security policy could block a user, an authentication policy could force the user to use MFA, or a decryption policy could force the user's traffic to be decrypted.



### Example Use Cases

Two DUG use case examples are shown here:





The first example shows how a user's entire known security state, which is derived from various sources, can determine how the firewall controls or affects the user's access to network resources. In this case, the user's network traffic is logged so that it can be analyzed. User metadata might also be collected from other resources, such as an LDAP server.

All this data can be analyzed in the firewall's logs, on Security information and event management (SIEM), in a user and entity behavior analytics system, or by using a variety of tools available at a security operations center (SOC). Any of these tools can be configured to tag or untag a username, depending on the results of the analysis. Tagging and untagging a username determines whether it is a member of a DUG. Then, DUG membership and policy configuration determine how the firewall should treat the user's network traffic.

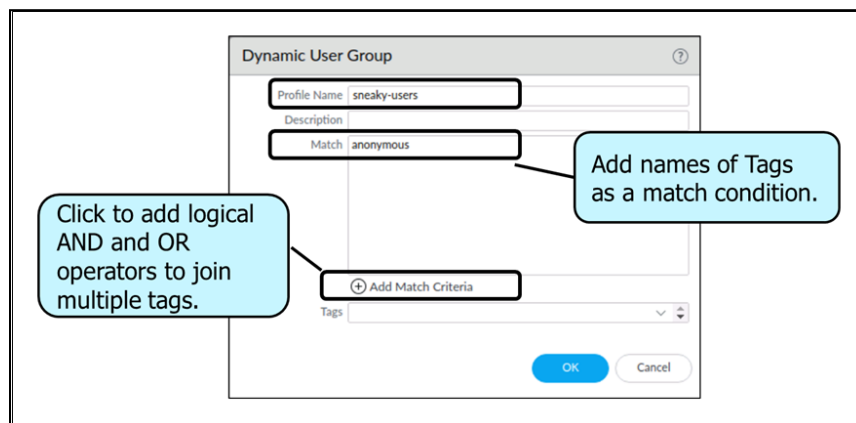
The second example illustrates how to use a DUG to implement time-based access controls for workers who might require only short-term access to network resources. In this case, you create a DUG and add it to policies that control user access to network resources. You then can add a time-based tag to a username. A tagged username is a member of the DUG, and network access is permitted by the DUG. When the time-based tag expires, the user's membership in the DUG is terminated, along with the network access that was provided by the DUG.

### 3.4.7 Requirements to support dynamic user groups

#### Configuring DUGs

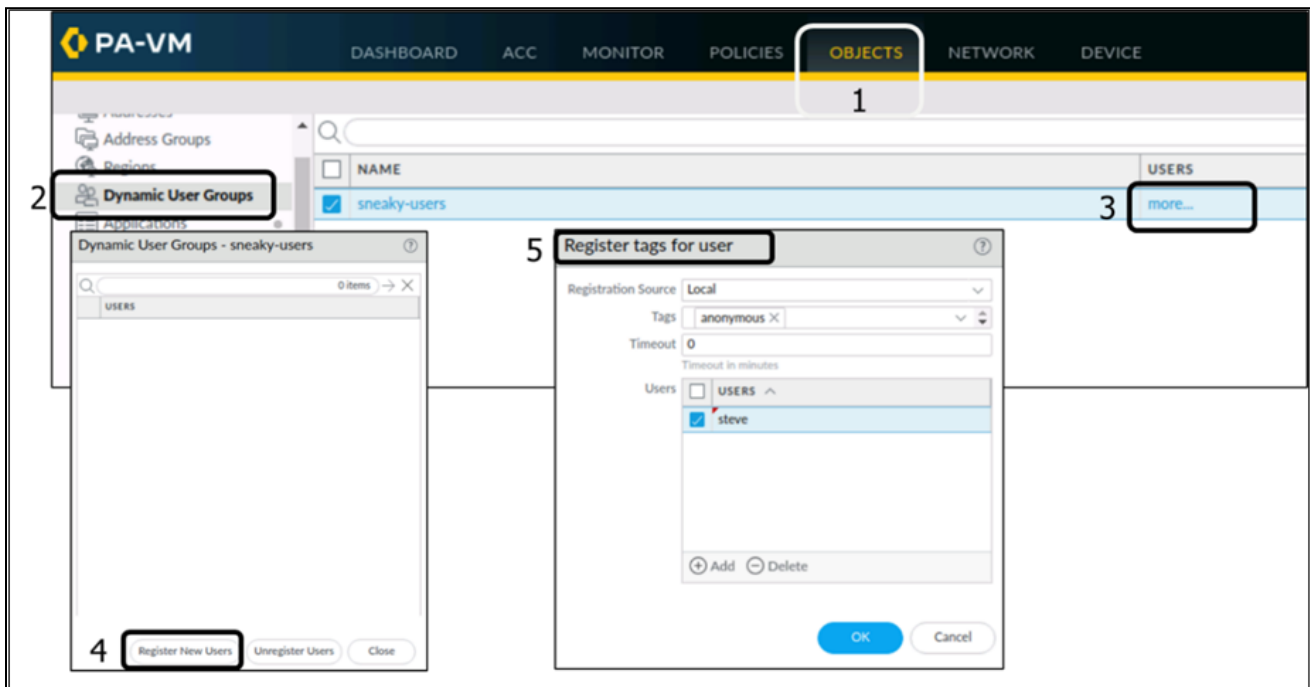
Before you can configure and use DUGs, first you need to configure User-ID in the environment. The User-ID agent is used to maintain the list of tags associated with usernames. Then, you configure the custom tags to use as match conditions for the DUG membership. In the web interface, browse to **Objects > Tags** and create one or more custom tags to assign dynamically to users. After creating the custom tags, create the DUGs.

To create and configure a DUG, browse to **Objects > Dynamic User Group** in the web interface. Then, click Add to create a new group. Provide a **Name** for the new group, optionally provide a **Description**, and then in the **Match** field, type one or more tags as match conditions. In the following example, the **Name** is **sneaky-users**, there is no **Description**, and the only **Match** condition is the tag named **anonymous**. If you click **Add Match Criteria**, you can use the logical AND and OR operators to join multiple tags as match conditions. The Tags field value is the tag that is statically assigned to the DUG object itself. It is not assigned to a user and not used as a match condition to identify users to add as DUG members.

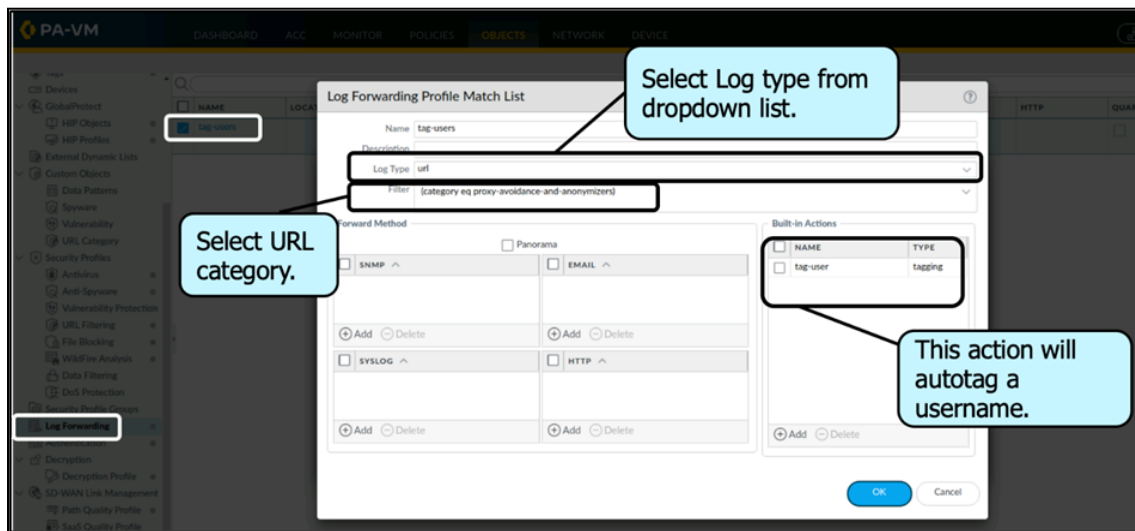


After creating the DUG, you configure the firewall to use the DUG. Four options are available. To dynamically register a tag with a username, you can use Panorama, the XML API, a remote User-ID agent, or the web interface. A firewall can forward the username and tag registration information to Panorama, and Panorama can distribute this information to the other firewalls. Other applications can invoke the firewall XML API commands to register the username and tag associations. A remote User-ID agent can forward the username and tag registrations to Panorama or other firewall User-ID agents. You also can use the web interface to register (or unregister) tags with usernames.

For example, the following screenshot illustrates the use of the web interface to register or unregister tags from a username. Browse to **Objects > Dynamic User Groups** and click **more** next to a group name. In the window that opens, click **Register New Users** to register a tag with a username. In the next window that opens, select the **Registration Source**. You can choose the local User-ID agent, a remote User-ID agent, or Panorama. In the following example, the **Local User-ID** agent was selected. Then, select the **Tags** to register with the user. The example uses the **anonymous** tag. If you want the tag to time out, which means the tag will be disassociated with the user, choose a **Timeout** value in minutes. Then, click **Add** and add one or more users to which to register the tag. To disassociate a tag from a username, start by clicking the **Unregister Users** button.



As a second example, you also can use a Log Forwarding Profile attached to a Security policy rule to auto-tag a username in response to a user's network behavior. In the following example, a Log Forwarding Profile named tag-users has been created. The profile is attached to a Security policy rule. If the rule matches an HTTP session and the URL Filtering log creates an entry where the URL category equals anonymous-proxy, then the Log Forwarding Profile invokes the built-in action that associates the tag anonymous with the username. The username is tagged and becomes a member of a DUG. Assuming that the DUG is used in the Security policy as a match condition, the firewall modifies what the user has access to.



### 3.4.8 How GlobalProtect internal and external gateways can be used

In a GlobalProtect mixed internal and external gateway configuration, you can set up two separate gateways for VPN access and for access to all sensitive internal resources. To do this, the GlobalProtect app performs internal host detection to determine if it is on the internal or external network.

### 3.4.9 References

- Mixed Internal and External Gateway Configuration, <https://docs.paloaltonetworks.com/globalprotect/10-1/globalprotect-admin/globalprotect-quick-configs/mixed-internal-and-external-gateway-configuration>
- User-ID Agent, <https://docs.paloaltonetworks.com/compatibility-matrix/user-id-agent>

## 3.5 Configure WildFire

The WildFire Analysis sandbox identifies previously unknown malware and generates signatures that the Palo Alto Networks firewalls can use to detect and block malware. When a Palo Alto Networks firewall is instructed to forward files and URLs via a WildFire Analysis Profile, the firewall can automatically forward the sample for WildFire analysis. WildFire determines the sample to be benign, grayware, phishing, or malicious based on the properties, behaviors, and activities that the

sample displays when analyzed and executed in the WildFire sandbox. WildFire then generates signatures to recognize the newly discovered malware and makes the latest signatures globally available every five minutes. Firewalls without a WildFire subscription license get the signature updates the following day, and firewalls with a WildFire license gain access to signatures within five minutes of generation. All the Palo Alto Networks firewalls worldwide then can compare incoming samples against these signatures to block the malware first detected by a single firewall automatically.

### 3.5.1 Submission profile

A WildFire submissions log is an automatically generated, time-stamped file that provides an audit trail to track events when a Palo Alto Networks network security platform forwards samples (files and emails links) to the WildFire cloud for analysis based on the WildFire Analysis profiles settings (**Objects > Security Profiles > WildFire Analysis**). The firewall generates the WildFire Submissions log entries for each sample it forwards after WildFire completes the static and dynamic analysis of the sample. WildFire Submissions log entries include the firewall Action for the sample (allow or block), the WildFire verdict for the submitted sample, and the severity level of the sample.

### 3.5.2 Action profile

WildFire Action settings in the Antivirus profile might impact traffic if the traffic generates a WildFire signature that results in a reset or a drop action. You can exclude internal traffic, such as software distribution applications through which you deploy custom-built programs, to transition safely (PAN-OS 9.1, 10.0, 10.1, 10.2) to best practices because WildFire might identify custom-built programs as malicious and generate a signature for them. Check **Monitor > Logs > WildFire Submissions** to see if any internal custom-built programs trigger WildFire signatures.

### 3.5.3 Submissions and verdicts

When WildFire analyzes a previously unknown sample in one of the Palo Alto Networks-hosted WildFire public clouds or a locally-hosted WildFire private cloud, a verdict is produced to identify samples as malicious, unwanted (grayware is considered obtrusive but not malicious), phishing, or benign. The following table summarizes the WildFire verdicts:

VERDICT	DESCRIPTION
Benign	Indicates that the entry received a WildFire analysis verdict of benign. Files categorized as benign are safe and do not exhibit malicious behavior.
Grayware	Indicates that the entry received a WildFire analysis verdict of grayware. Files categorized as grayware do not pose a direct security threat but might display otherwise obtrusive behavior. Grayware can include adware, spyware, and Browser Helper Objects (BHOs).
Phishing	Indicates that WildFire assigned a link and analysis verdict of phishing. A phishing verdict indicates that the site to which the link directs users displayed credential phishing activity.
Malicious	Indicates that the entry received a WildFire analysis verdict of malicious. Samples categorized as malicious can pose a security threat. Malware can include viruses, C2 (command-and-control), worms, Trojans, Remote Access Tools (RATs), rootkits, and botnets. For samples that are identified as malware, the WildFire cloud generates and distributes a signature to prevent against future exposure.

Each WildFire cloud—global (U.S.), regional, and private—analyzes samples and generates WildFire verdicts independently of the other WildFire clouds. With the exception of WildFire private cloud verdicts, WildFire verdicts are shared globally, enabling WildFire users to access a worldwide database of threat data.

### 3.5.4 Signature actions

WildFire can discover zero-day malware in web traffic (HTTP/HTTPS), email protocols (SMTP, IMAP, and POP), and FTP traffic and can quickly generate signatures to identify and protect against future infections from the malware it discovers. WildFire automatically generates a signature based on the malware payload of the sample and tests it for accuracy and safety.

Each WildFire cloud—global, regional, and private—analyzes samples and generates malware signatures independently of the other WildFire clouds. With the exception of WildFire private cloud signatures, WildFire signatures are shared globally, enabling WildFire users worldwide to benefit from malware coverage regardless of the location in which the malware was first detected. Because malware evolves rapidly, the signatures that WildFire generates address multiple variants of the malware.

Firewalls with an active WildFire license can retrieve the latest WildFire signatures in real-time, as soon as they become available. If you do not have a WildFire subscription, signatures are made available within 24-48 hours as part of the antivirus update for firewalls with an active Threat Prevention license.

As soon as the firewall downloads and installs the new signature, the firewall can block the files that contain that malware (or a variant of the malware). Malware signatures do not detect malicious and phishing links; to enforce these links, you must have a PAN-DB URL Filtering license. You can then block user access to malicious and phishing sites.

### 3.5.5 File types and file sizes

#### File Types

The following table lists the supported file categories. The example file types represent commonly used formats rather than a complete list; this list changes as the underlying technology supports new file types. Use the examples to gauge the scope of the file categories.

FILE TYPE CATEGORY	WILDFIRE ANALYSIS SUPPORT	EXAMPLE FILE TYPES
Android application package	✓	apk, ETC.
Adobe Flash	✓	.swf, ETC.
Java Archive	✓	jar, ETC.
Microsoft Office	✓	docx, xlsx, pptx, ooxml, ETC.
Portable executable	✓	pe, exe, ETC.
Portable Document Format	✓	pdf, ETC.
Mac OS X	✓	dmg, pkg, ETC.
Archive	✓	rar, 7z, ETC.
Linux	✓	elf, ETC.
Script	✓	bat, js, vbs, ps1, ETC.

## File Sizes

The maximum and default WildFire file forwarding sizes and rates are increased in PAN-OS® 9.0 to provide optimal visibility and detection. Based on the Palo Alto Network's data analytics, the new default capacities protect against the majority of threats and best practice is to use the new default values.

FILE TYPE	PAN-OS 9.0 DEFAULT FILE FORWARDING SIZES	PAN-OS 9.0 SIZE LIMITS
pe	16MB	1-50MB
apk	10MB	1-50MB
pdf	3,072KB	100-51,200KB
ms-office	16,384KB	200-51,200KB
jar	5MB	1-20MB
flash	5MB	1-10MB
MacOSX	10MB	1-50MB
archive	50MB	1-50MB
linux	50MB	1-50MB

### 3.5.6 Update schedule

The Palo Alto Networks NGFW supports the real-time retrieval of WildFire® signatures. This enables you to access the signatures as soon as they are generated, which greatly minimizes the window in which malware can infiltrate a network. Signature downloads that occur during a sample check are saved in the firewall cache and are available for fast (local) look-ups. In addition, to maximize coverage, the firewall automatically downloads a supplementary signature package on a regular basis when you enable real-time signatures. These signatures remain available in the firewall cache until they become stale and are refreshed or are overwritten by new signature updates. Palo Alto Networks determines which protections are the most relevant and timely and includes those in the signature packages.

### 3.5.7 Forwarding of decrypted traffic

Enable the firewall to forward decrypted SSL traffic for Advanced WildFire analysis. Traffic that the firewall decrypts is evaluated against Security policy rules; if it matches the WildFire analysis profile attached to the security rule, the decrypted traffic is forwarded for analysis before the firewall re-encrypts it. Only a super user can enable this option.

On a firewall that does not have multiple virtual systems enabled:

- If not done already, enable the firewall to perform decryption and Forward Files for Advanced WildFire Analysis.
- Select **Device > Setup > Content-ID**.
- Edit the Content-ID settings and **Allow Forwarding of Decrypted Content**.
- Click **OK** to save the changes.

On a firewall with virtual systems enabled:

- If not done already, enable decryption and Forward Files for Advanced WildFire Analysis.
- Select **Device > Virtual Systems**, click the virtual system you want to modify, and **Allow Forwarding of Decrypted Content**.

### 3.5.8 References

- WildFire Submissions Logs,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/view-and-manage-logs/log-types-and-severity-levels/wildfire-submissions-logs>
- WildFire Signatures,  
<https://docs.paloaltonetworks.com/wildfire/10-2/wildfire-admin/wildfire-overview/wildfire-concepts/wildfire-signatures>
- Supported File Types for WildFire Analysis,  
<https://docs.paloaltonetworks.com/saas-security/saas-security-admin/saas-security-api/get-started-with-saas-security-api/support-on-saas-security-api/supported-file-types-for-wildfire-analysis>
- Increased WildFire File Forwarding Capacity,  
<https://docs.paloaltonetworks.com/wildfire/u-v/wildfire-whats-new/wildfire-features-in-panos-90/increased-wildfire-file-forwarding-capacity>
- WildFire Real-Time Signature Updates,  
<https://docs.paloaltonetworks.com/wildfire/u-v/wildfire-whats-new/wildfire-features-in-panos-100/wildfire-real-time-signature-updates>
- Forward Decrypted SSL Traffic for Advanced WildFire Analysis,  
<https://docs.paloaltonetworks.com/advanced-wildfire/administration/configure-advanced-wildfire-analysis/forward-decrypted-ssl-traffic-for-advanced-wildfire-analysis>

## 3.6 Configure Web Proxy

If your network uses a proxy device for security, you can leverage the same level of protection by using the on-premises web proxy capability with PAN-OS 11.0. The web proxy feature enables additional options for migrating from an existing web proxy architecture to a simple unified management console. Using the web proxy feature with Prisma Access provides a seamless method for migrating, deploying, and maintaining secure web gateway (SWG) configurations from an easy-to-use and simplified interface. Web proxy helps during the transition from on-premises to the cloud with no loss to security or efficiency. Web proxy requires both a valid DNS Security license and the Prisma Access explicit proxy license.

The web proxy supports two methods for routing traffic:

- Explicit Proxy
- Transparent Proxy

The following platforms support web proxy:

- PA-1400
- PA-3400
- VM Series (with vCPUs)
- Panorama using PAN-OS 11.0

### 3.6.1 Transparent proxy

For the transparent proxy method, the request contains the destination IP address of the web server and the client browser is redirected to the proxy. There is no client configuration and Panorama is optional. Transparent proxy requires a loopback interface, User-ID configuration in the proxy zone, and specific Destination NAT (DNAT) rules. Transparent proxy does not support X-Authenticated-User (XAU).

### 3.6.2 Explicit proxy

For the explicit proxy method, the request contains the destination IP address of the configured proxy and the client browser sends requests to the proxy directly.

### 3.6.3 References

- Configure a Web Proxy,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/dns/configure-a-web-proxy>



## Domain 4: Deploy and Configure Firewalls Using Panorama

### 4.1 Configure templates and template stacks

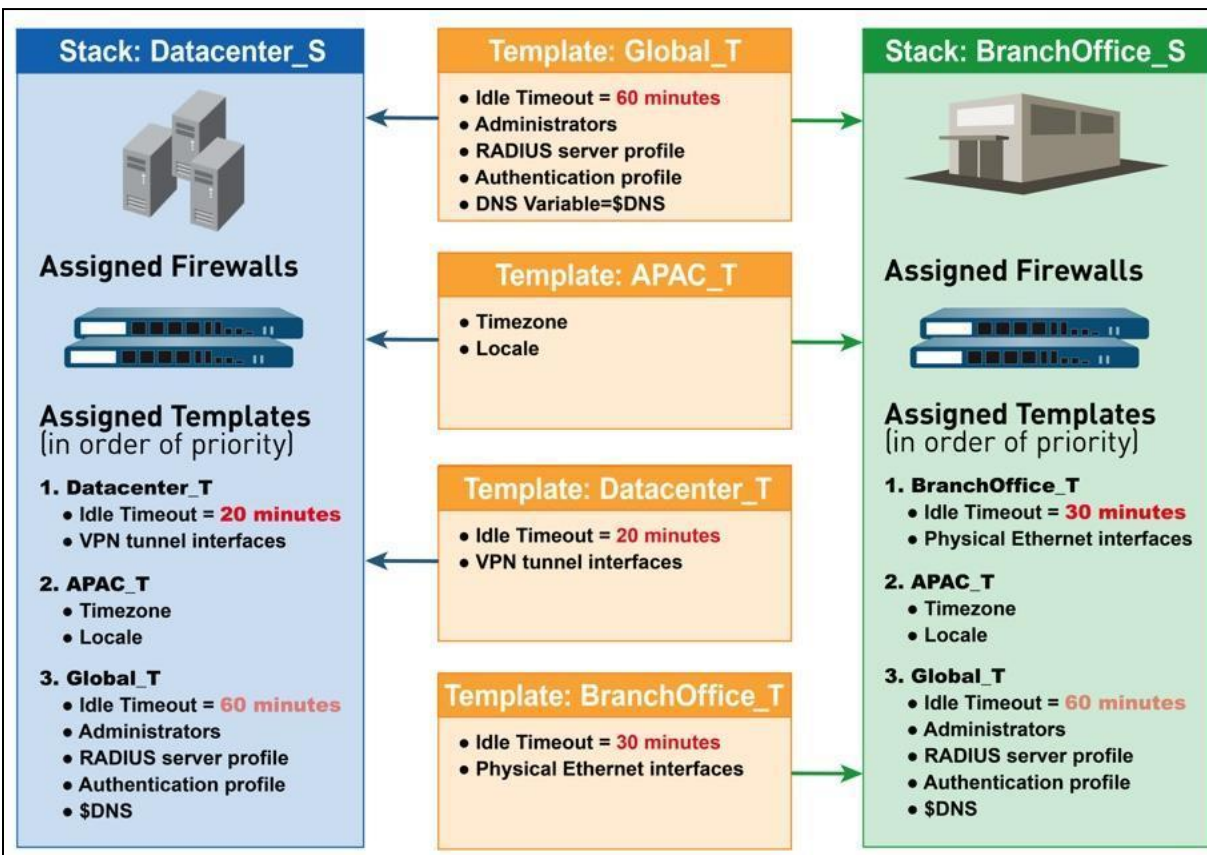
For centralized firewall configuration and update management, define a template stack.

#### 4.1.1 Components configured in a template

To create a template stack, navigate to **Panorama > Templates > Add Stack** and then name the stack. Template stacks can combine up to eight templates. Add templates in the order of priority. Next, in the **Device** section, select firewalls to assign them to the stack. You can assign any firewall to only one template stack. Optionally, select **Group HA Peers to** for the firewalls in HA configuration.

#### 4.1.2 How the order of templates in a stack affects the configuration push to a firewall

Templates in a stack have a configurable priority order to ensure that Panorama pushes only one value for any duplicate setting. The following illustration shows a data center stack in which the data center template has a higher priority than the global template.



### 4.1.3 Overriding a template value in a stack

While templates and template stacks enable you to apply a base configuration to multiple firewalls, you might want to configure firewall-specific settings that don't apply to all the firewalls in a template or template stack. Conversely, you may want to override the template settings to create a template stack configuration that you can apply as a base configuration to all the managed firewalls. Overrides allow for exceptions or modifications to meet configuration needs. For example, if you use a template to create a base configuration but a few firewalls in a test lab environment need different settings for the DNS server IP address or the Network Time Protocol server, you can override the template and template stack settings.

You can override a template or template stack value in one of the following ways:

- Using variables: Define a value locally on the firewall to override a value pushed from a template or template stack or define firewall-specific variables to override values pushed from a template or template stack.
- Using a template stack: Define values or variables on the template stack to override values pushed from a template.

### 4.1.4 Configure variables in templates

You can use template stack variables to replace IP addresses, group IDs, and interfaces in the configurations. Variables allow you to reduce the total number of templates and template stacks. This lets you use fewer templates and template stacks while using specific values that otherwise would have needed their own template or template stack.

### 4.1.5 Relationship between Panorama and devices for dynamic update versions, policy implementation, and HA peers

The firewall retrieves updates and uses them to enforce policy, without requiring configuration changes. You can view the latest updates, read the release notes for a description of each update, and then select the update you want to download and install.

#### 4.1.6 References

- Configure a Template or Template Stack Variable,  
<https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/manage-firewalls/manage-templates-and-template-stacks/configure-template-or-template-stack-variables>
- Templates and Template Stacks,  
<https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/panorama-overview/centralized-firewall-configuration-and-update-management/templates-and-template-stacks>
- Manage Templates and Template Stacks,  
<https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/manage-firewalls/manage-templates-and-template-stacks>
- Template Capabilities and Exceptions,  
<https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/manage-firewalls/manage-templates-and-template-stacks/template-capabilities-and-exceptions>
- Device > Dynamic Updates,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/device/device-dynamic-updates>
- Override a Template or Template Stack Value,  
<https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/manage-firewalls/manage-templates-and-template-stacks/override-a-template-setting>

## 4.2 Configure device groups

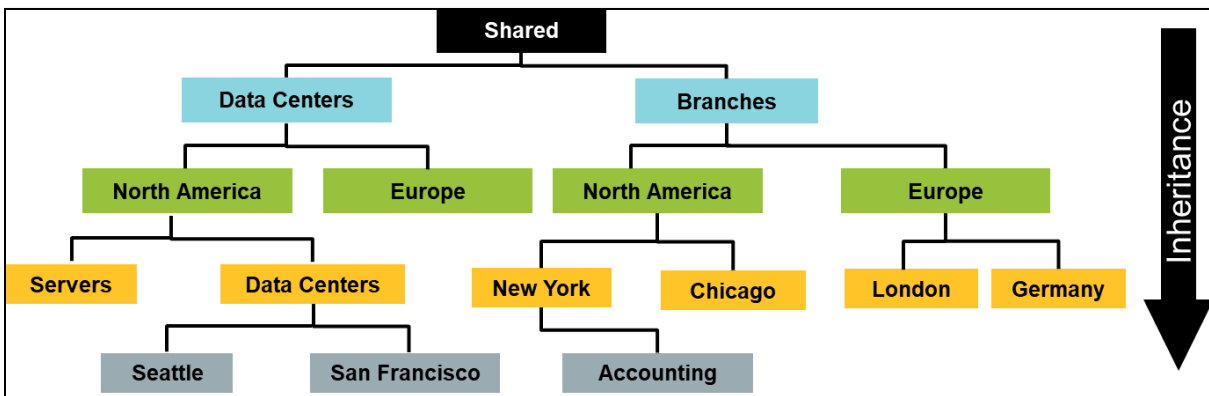
Before using Panorama effectively, you must group the firewalls in the network into logical units called device groups. A device group enables grouping based on network segmentation, geographic location, organizational function, or any other common aspect of firewalls that require similar policy configurations. You can use device groups to configure policy rules and the objects they reference. You can organize device groups hierarchically, with shared rules and objects at the top and device group-specific rules and objects at subsequent levels. Organization enables you to create a hierarchy of rules that enforces how firewalls handle traffic. For example, you can define a set of shared rules as a corporate acceptable use policy. Then, to allow only regional offices to access peer-to-peer traffic, such as BitTorrent, you can define a device group rule that Panorama pushes only to the regional offices (or define a shared Security rule and target it to the regional offices).

### 4.2.1 Device group hierarchies

#### Device Groups

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

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



#### 4.2.2 Identify what device groups contain

Device groups enable a layered approach for managing policies across a network of managed firewalls. A firewall evaluates policy rules by layer (shared, device group, and local) and by type (pre-rules, post-rules, and default rules) in the order shown in the following figure:

Name	Tags	Type	Source				Destination		Rule Usage		
			Zone	Address	User	HIP Profile	Zone	Address	Hit Count	Last Hit	First Hit
Allow Web	none	universal	Trust-L3	any	any	any	Untrust-L3	any	2285	2017-11-14 20:17:53	2017-11-11 21:52:58
Outbound FTP	none	universal	Trust-L3	any	any	any	Untrust-L3	any	0	-	-
Local Policy	none	universal	Trust-L3	any	any	any	DMZ	any	-	-	-
Allow Facebook	none	universal	Trust-L3	any	any	any	Untrust-L3	any	0	-	-
intrazone-default	none	intrazone	any	any	any	any	(intrazone)	any	6772	2017-11-14 20:17:23	2017-10-31 20:58:56
interzone-default	none	interzone	any	any	any	any	any	any	298702	2017-11-14 19:54:12	2017-10-31 21:02:50

- Pre-Policy Rules
- Local Policy Rules
- Post-Policy Rules
- Default Rules

#### 4.2.3 Differentiate between different use cases for pre-rules, local rules, default rules, and post-rules

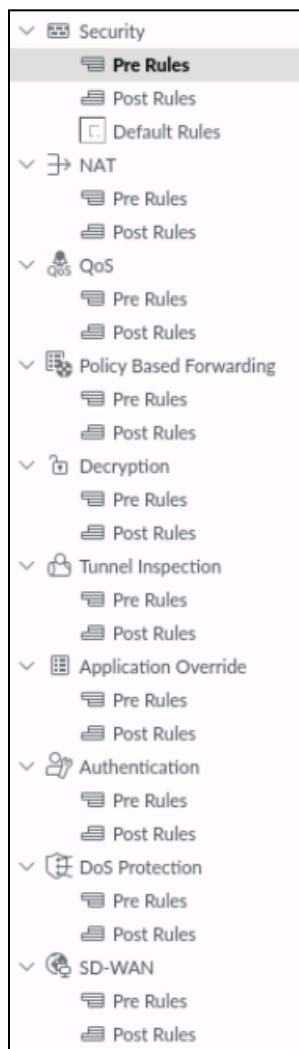
When the firewall receives traffic, it performs the action defined in the first evaluated rule that matches the traffic and disregards all the subsequent rules. Whether you view rules on a firewall or in Panorama, the web interface displays them in evaluation order. All the shared, device group, and default rules that the firewall inherits from Panorama are shaded orange. Local firewall rules display between the pre-rules and post-rules.

Objects are configuration elements that policy rules reference—for example, IP addresses, URL categories, Security Profiles, users, services, and applications. Rules of any type (pre-rules, post-rules, default rules, and rules locally defined on a firewall) and any rulebase (security, NAT, QoS, PBF,

decryption, Application Override, Captive Portal, and DoS protection) can reference objects. You can reuse an object in any number of rules that have the same scope as that object in the device group hierarchy.

For example, if you add an object to the shared location, all the rules in the hierarchy can reference that shared object because all the device groups inherit objects from the shared location. If you add an object to a particular device group, only the rules in that device group and its descendant device groups can reference that device group object. If object values in a device group must differ from those inherited from an ancestor device group, you can override the inherited object values. You can also revert to the inherited object values at any time. When you create objects for use in a shared or device group policy once and use them many times, you reduce the administrative overhead and ensure consistency across firewall policies.

When new policy rules are entered into a Panorama device group, the device group and the pre- or post-designation must be decided. The pre- and post-designations are chosen through the selection of the appropriate policy menu item, as shown in the following figure:



#### 4.2.4 Identify the impact of configuring a primary device

Every firewall and Panorama management server has a default master key that encrypts all the private keys and passwords in the configuration to secure them (such as the private key used for SSL Forward Proxy Decryption).

In an HA configuration, you must use the same master key on both firewalls because the master key is not synchronized across HA peers. Otherwise, HA synchronization will not work properly.

If you are using Panorama to manage firewalls, you can configure the same master key on Panorama and all the managed firewalls or configure a unique master key for each managed firewall. For managed firewalls in an HA configuration, you must configure the same master key for each HA peer.

Be sure to store the master key in a safe location. You cannot recover the master key, and the only way to restore the default master key is to reset the firewall to factory default settings.

#### 4.2.5 Assign firewalls to device groups

Device groups comprise firewalls and virtual systems you want to manage as a group, such as the firewalls that manage a group of branch offices or individual departments in a company. Panorama treats these groups as single units when applying policies. Firewalls can belong to only one device group, but because virtual systems are distinct entities in Panorama, you can assign a virtual system within a firewall to different device groups.

You can nest device groups in a tree hierarchy of up to four levels under a shared location to implement a layered approach for managing policies across the network of firewalls. At the bottom level, a device group can have parent, grandparent, and great-grandparent device groups at successively higher levels—collectively called ancestors—from which the bottom-level device group inherits policies and objects. At the top level, a device group can have child, grandchild, and great-grandchild device groups—collectively called descendants. When you select **Panorama > Device Groups**, the **Name** column displays this device group hierarchy.

After adding, editing, or deleting a device group, perform a Panorama commit and device group commit. Panorama then pushes the configuration changes to the firewalls that are assigned to the device group. Panorama supports up to 1,024 device groups.

#### 4.2.6 References

- Configure the Master Key, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/certificate-management/configure-the-master-key>

### 4.3 Manage firewall configurations within Panorama

#### 4.3.1 Licensing

The Panorama Software Firewall License plugin allows you to automatically license a VM-Series firewall when it connects to Panorama. If the VM-Series firewalls are located in the perimeter of

deployment and do not have connectivity to the Palo Alto Networks licensing server, the Software Firewall License plugin simplifies the license activation process by using Panorama to license the VM-Series firewall.

Additionally, the Software Firewall License plugin simplifies the license activation and deactivation of VM-Series firewalls in environments that use auto-scaling and automation to deploy and delete firewalls to address changes in the cloud.

To install the Panorama Software Firewall License plugin, you must use Panorama 10.0.0 or later and VM-Series plugin 2.0.4 or later. Your VM-Series firewalls must run PAN-OS 9.1.0 or later.

### 4.3.2 Commit recovery feature

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

### 4.3.3 Automatic commit recovery

Panorama automatic commit recovery enables you to configure the firewall to attempt a specified number of connectivity tests after you push a configuration change from Panorama or commit a configuration change locally on the firewall. Automatic commit recovery is enabled by default, thus enabling the managed firewalls to locally test the configuration pushed from Panorama to verify that the new changes do not break the connection between Panorama and the managed firewall. If the committed configuration breaks the connection, then the firewall automatically fails the commit and the configuration reverts to the previous running configuration.

The firewall also checks the connectivity to Panorama every hour to ensure consistent communication if unrelated network configuration changes have disrupted connectivity between the firewall and Panorama or if implications to a pushed committed configuration might have affected connectivity. If an hourly connectivity check fails, the firewall generates a system log to alert administrators of potential configuration or network connectivity issues. An event is generated in the system log when you disable the setting, when a connectivity test fails, or when a firewall configuration reverts to the last running configuration.

In HA firewall configurations, each HA peer performs the connectivity tests independently of each other. HA configuration syncs might occur only after each HA successfully tests the connectivity to Panorama and verifies its connection.

### Configuration Settings for Panorama Automatic Commit Recovery

PAN-OS allows the managed firewalls to check for connectivity to the Panorama management server and to revert automatically to the last running configuration when the firewall is unable to communicate with Panorama.

Automatic commit recovery enables you to configure the firewall to attempt a specified number of connectivity tests, as well as the interval at which each test occurs. Connectivity tests occur before

the firewall automatically reverts its configuration to the previous running configuration after you push a configuration from Panorama or locally commit a configuration change on the firewall.

#### 4.3.4 Commit types and schedules

To commit types and schedules:

1. Log in to the Panorama web interface.
2. Select **Device > Setup > Management**.
3. In the **Template context** drop-down list, select the template or template stack that manages the devices for which you want to configure the automated commit recovery parameters. Configure the automated commit recovery settings:
  - a. Edit (⚙️) the **Panorama Settings**.
  - b. Verify that **Enable automated commit recovery** is selected.
  - c. Enter a value in the **Number of attempts to check for Panorama connectivity** field.
  - d. Enter a value in the **Interval between retries** field.
  - e. Click **OK** to save the configuration changes.
4. Select **Commit** and **Commit and Push** the configuration changes.

Panorama Settings

Panorama Servers

192.168.1.252

Enable pushing device monitoring data to Panorama

Receive Timeout for Connection to Panorama (sec) 240

Send Timeout for Connection to Panorama (sec) 240

Retry Count for SSL Send to Panorama 25

Enable automated commit recovery

Number of attempts to check for Panorama connectivity 1

Interval between retries (sec) 10

Disable Panorama Policy and Objects Disable Device and Network Template OK Cancel

Verify that the automated commit recovery feature is enabled on the managed firewalls.

5. Launch the firewall web interface.
6. Select **Device > Setup > Management**. In **Panorama Settings**, verify that **Enable automated commit recovery** is selected.



### 4.3.5 Configuration backups

#### Running Configuration and Candidate Configuration

Firewall settings are stored in the XML configuration files that can be archived, restored, and managed. A firewall contains both a running configuration that contains all of the settings that are currently active and a candidate configuration. The candidate configuration is a copy of the running configuration, which also includes any settings changes that are not yet committed. Changes made using the management web interface, the CLI, or the XML API are staged in the candidate configuration until you perform a commit operation. During a commit operation, the candidate configuration replaces the running configuration.

#### Panorama and Firewall Configuration Backups and Restorations

When Panorama has a management relationship with a firewall, Panorama can obtain copies of both Panorama-managed and locally managed configurations. After a commit on a local firewall that runs PAN-OS 5.0 or later, a backup is sent of the running configuration to Panorama. Any commits performed on the local firewall triggers the backup, including any commits that an administrator performs locally on the firewall or that PAN-OS initiates and automatically commits (such as an FQDN refresh).

By default, Panorama stores up to 100 backups for each firewall although this is configurable. To store Panorama and firewall configuration backups on an external host, you can schedule exports from Panorama or complete an export on demand. The saved configuration files can be restored to the firewall at any time by an administrator by using the **Panorama > Managed Devices > Summary** tools.

#### Return Merchandise Authorization Replacement of a Panorama-Managed Firewall

To minimize the effort required to restore the configuration on a managed firewall, you can use a Return Merchandise Authorization to replace the serial number of the old firewall with that of the new firewall on Panorama. To then restore the configuration on the replacement firewall, either import a firewall state that was previously generated and exported from the firewall or use Panorama to generate a partial device state for the managed firewalls running PAN-OS 5.0 and later versions. By replacing the serial number and importing the firewall state, you can resume using Panorama to manage the firewall.

### 4.3.6 Commit type options

Click Commit at the top right of the web interface and select an operation for pending changes to the Panorama configuration and changes that Panorama pushes to firewalls, Log Collectors, and WildFire clusters and appliances:

- **Commit > Commit to Panorama** — Activates the changes you made in the configuration of the Panorama management server. This action also commits device group, template, Collector Group, and WildFire cluster and appliance changes to the Panorama configuration without pushing the changes to firewalls, Log Collectors, or WildFire clusters and appliances. Committing just to the Panorama configuration enables you to save the changes that are not ready for activation on the firewalls, Log Collectors, or WildFire clusters and appliances.

- **Commit > Push to Devices** — Pushes the Panorama running configuration to device groups, templates, Collector Groups, and WildFire clusters and appliances.
- **Commit > Commit and Push** — Commits all the configuration changes to the local Panorama configuration and then pushes the Panorama running configuration to device groups, templates, Collector Groups, and WildFire clusters and appliances.

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

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

The following options are available for committing, validating, or previewing configuration changes:

- The following options apply when you commit to Panorama by selecting **Commit > Commit to Panorama** or **Commit > Commit and Push**:
  - Commit All Changes
  - Commit Changes Made By
  - Commit Scope
  - Location Type
  - Object Type
  - Admins
  - Include in Commit
  - Group by Type
  - Preview Changes
  - Change Summary
  - Validate Commit

- The following options apply when you push configuration changes to managed devices by selecting **Commit > Push to Devices** or **Commit > Commit and Push**:
  - Push All Changes
  - Push Changes Made By
  - Push Scope
  - Location Type
  - Object Type
  - Entities
  - Admins
  - Include in Push
  - Edit Selections
  - Device Groups and Templates
  - Filters
  - Name
  - Last Commit State
  - HA Status
  - Changes Pending (Panorama) Commit
  - Preview Changes column
  - Select All
  - Deselect All
  - Expand All
  - Collapse All
  - Group HA Peers
  - Validate
  - Filter Selected
  - Merge with Candidate config
  - Include Device and Network Templates
  - Force Template Values
  - Log Collector Groups
  - WildFire Appliances and Clusters
  - Filters
  - Name
  - Last Commit State
  - No Default Selections
  - Validate Device Group Push
  - Validate Template Push
  - Group by Location Type
  
- The following options apply when you commit the Panorama configuration or push changes to devices:
  - Description
  - Commit / Push / Commit and Push

### 4.3.7 Manage dynamic updates for Panorama and Panorama-managed devices

Setting a schedule for dynamic updates allows you to define the frequency at which the firewall checks for and downloads or installs new updates. Particularly for Applications and Threats content updates, you might want to set a schedule that staggers new and modified application updates behind the threat updates; this gives you more time to assess how new and modified applications impact your Security policy while ensuring that the firewall is always equipped with the latest threat protections.

Panorama requires a direct Internet connection for scheduling Supported Updates on firewalls, Log Collectors, and WildFire appliances and appliance clusters. Otherwise, you can perform only on-demand updates. (To schedule Antivirus, WildFire, or BrightCloud URL updates for Log Collectors, the Log Collectors must be running Panorama 7.0.3 or a later release.) Each firewall, Log Collector, or WildFire appliance or appliance cluster receiving an update generates a log to indicate that the installation succeeded (a Config log) or failed (a System log).

### 4.3.8 Software and dynamic updates

#### Dynamic Updates

Palo Alto Networks frequently publishes dynamic updates to the firewall. This allows for security updates without the need to upgrade firmware.

#### Software Updates

To ensure that you are always protected from the latest threats (including those that have not yet been discovered), you must ensure to keep the firewalls up-to-date with the latest content and software updates published by Palo Alto Networks. The Dynamic Content Updates available depends on the subscriptions you have. You can set a schedule for content updates and define the frequency at which the firewall retrieves and installs updates.

### 4.3.9 Import firewall configurations into Panorama

If you have already deployed the Palo Alto Networks firewalls and locally configured them but now want to use Panorama to centrally manage them, you must perform pre-migration planning. This involves importing firewall configurations into Panorama and verifying that the firewalls function as expected after the transition. If some settings are unique to individual firewalls, you can continue accessing the firewalls to manage the unique settings. You can manage any firewall setting by pushing its value from Panorama or by configuring it locally on the firewall, but you cannot manage the setting through both Panorama and the firewall. If you want to exclude certain firewall settings from Panorama management, you can either:

- Migrate the entire firewall configuration and then, on Panorama, delete the settings that you will manage locally on firewalls. You can override a template or template stack value that Panorama pushes to a firewall instead of deleting the setting on Panorama.
- Load a partial firewall configuration, including only the settings that you will use Panorama to manage.

Firewalls do not lose logs during the transition to Panorama management.

### 4.3.10 Configure Log Collectors

Select **Panorama > Managed Collectors** to manage Log Collectors. When you add a new Log Collector as a managed collector, the settings you configure vary based on the location of the Log Collector and whether you deployed Panorama in an HA configuration. The settings include:

- **Dedicated Log Collector:** The **Interfaces** tab is not initially displayed when you add a Log Collector. You must enter the serial number (**Collector S/N**) of the Log Collector, click **OK**, and then edit the Log Collector to display the interface settings.
- **Default Log Collector local to the solitary (non-HA) or active (HA) Panorama management server:** After you enter the serial number (**Collector S/N**) of the Panorama management server, the **Collector** dialog displays only the disks, communication settings, and a subset of the general settings. The Log Collector derives its values for all the other settings from the configuration of the Panorama management server.
- **(HA only) Default Log Collector local to the passive Panorama management server:** Panorama treats this Log Collector as remote, so you must configure it as you would configure a dedicated Log Collector.

### 4.3.11 Check firewall health and status from Panorama

Panorama allows you to monitor the hardware resources and performance for managed firewalls. Panorama centralizes time-trended performance information (CPU, memory, CPS, and throughput) and logging performance, environmental information (fans, RAID status, and power supplies), and correlates events—such as commits, content installs, and software upgrades—to health data. When a firewall deviates from its calculated baseline, Panorama reports it as a deviating device to help identify, diagnose, and resolve any hardware issues quickly.

### 4.3.12 Configure role-based access control on Panorama

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

### 4.3.13 References

- For more information, refer to the link below, <https://docs.paloaltonetworks.com/vm-series/11-0/vm-series-deployment/license-the-vm-series-firewall/use-panorama-based-software-firewall-license-management>
- Panorama Commit, Validation, and Preview Operations, <https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/panorama-overview/panorama-commit-validation-and-preview-operations>
- Enable Automated Commit Recovery, <https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/administer-panorama/enable-automated-commit-recovery>
- Schedule Dynamic Content Updates, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/panorama-web-interface/panorama-device-deployment/schedule-dynamic-content-updates>
- Install Content Updates, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-upgrade/software-and-content-updates/install-content-and-software-updates>
- Manage Configuration Backups, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/firewall-administration/manage-configuration-backups>
- Manage Panorama and Firewall Configuration Backups, <https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/administer-panorama/manage-panorama-and-firewall-configuration-backups>
- Panorama Commit Operations, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/panorama-web-interface/panorama-commit-operations>
- Device > Dynamic Updates, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/device/device-dynamic-updates>
- Schedule a Content Update Using Panorama, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-upgrade/upgrade-panorama/deploy-updates-to-firewalls-log-collectors-and-wildfire-appliances-using-panorama/schedule-a-content-update-using-panorama>
- Replace an RMA Firewall, <https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/troubleshooting/replace-an-rma-firewall>
- Backing up and Restoring Configurations, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIRcCAK>
- Panorama > Managed Devices > Health, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/panorama-web-interface/panorama-managed-devices-summary/panorama-managed-devices-health>
- Transition a Firewall to Panorama Management, <https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/manage-firewalls/transition-a-firewall-to-panorama-management.html>
- Log Collector Configuration, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/panorama-web-interface/panorama-managed-collectors/log-collector-configuration>

## Domain 5: Manage and Operate

### 5.1 Manage and configure log forwarding

#### 5.1.1 Identify log types and criticalities

##### Log forwarding, Filtering, and Tagging

Log Forwarding Profiles can be used to filter and forward logs from the following firewall logs:

- Authentication
- Data Filtering
- Decryption
- Traffic
- Threat
- Tunnel
- URL Filtering
- WildFire Submissions

##### Methods Used to Forward Logs

Depending on the log message type, two main methods are used to forward log events: redirecting log events based on event types and redirecting log events to different systems.

Log events destined for the System, Config, User-ID, HIP Match, and IP-Tag logs can be redirected using specific event types. These types can be configured in **Device > Log Settings**, as shown in the following figure:

**Method 1**

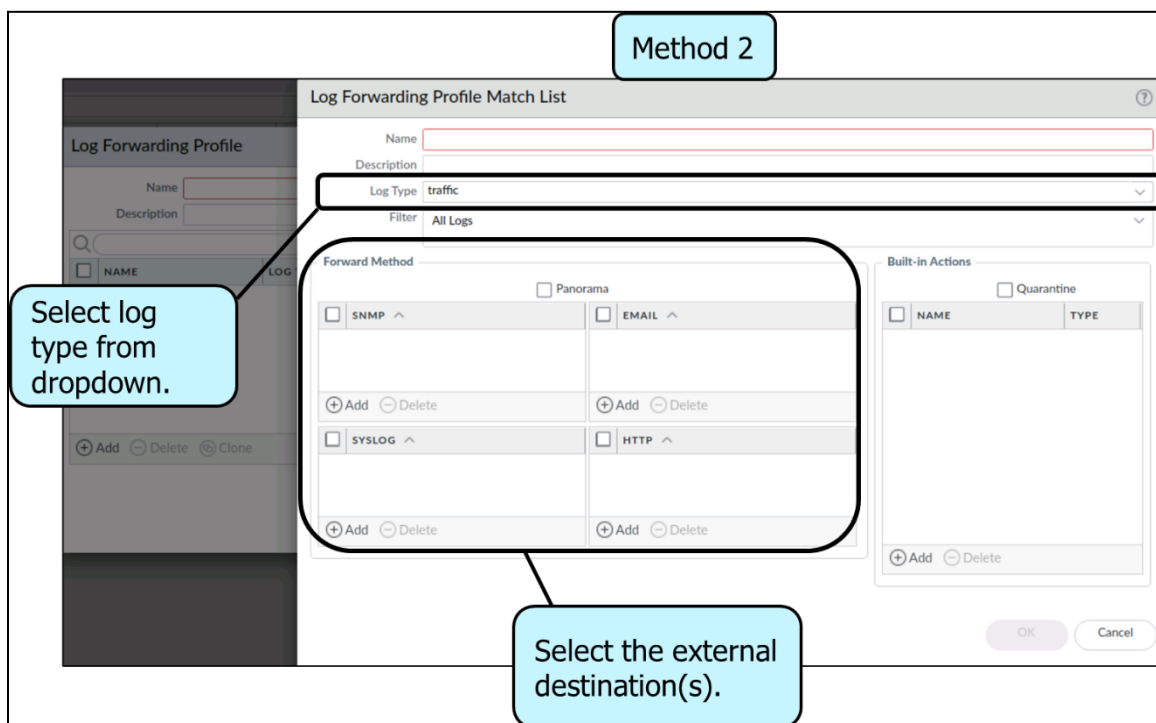
NAME	DESCRIPTION	FILTER	PANORAMA	SNMP TRAP
<input type="checkbox"/> system-informational		(severity eq informational)	<input checked="" type="checkbox"/>	
<input type="checkbox"/> system-low		(severity eq low)	<input checked="" type="checkbox"/>	
<input type="checkbox"/> system-medium		(severity eq medium)	<input checked="" type="checkbox"/>	
<input type="checkbox"/> system-high		(severity eq high)	<input checked="" type="checkbox"/>	

NAME	DESCRIPTION	FILTER	PANORAMA	SNMP TRAP	EMAIL
<input type="checkbox"/> config-any		All Logs	<input checked="" type="checkbox"/>		

Choose event destination(s) for specific event types.

Log events can also be redirected to other systems, such as Panorama, SIEM products, and the syslog server, using a Log Forwarding Profile. A Log Forwarding Profile can route traffic, threat, WildFire, and other log events, as shown in the following figure:



Log Forwarding Profiles are attached to individual firewall Security policy rules to enable forwarding events associated with specific policies. These profiles include one or more Log Forwarding Profile match lists. This granularity allows administrators specific control of forwarding and the potential to customize forwarding for policies of differing importance.

All forwarded events are sent to their destination as they are generated on the firewall. Palo Alto Networks also offers Cortex Data Lake, a cloud-based solution that can serve as a central repository for forwarded logs from multiple Palo Alto Networks devices. This central pool of log data is fully accessible to the owner, and it acts as an optional base for further third-party security applications through the Palo Alto Networks Cortex API.

### Log Forwarding Profiles

To maximize the efficiency of the incident response and monitoring operations, you can create custom log forwarding filters based on any log attributes, such as threat type or source user. Instead of forwarding all the logs or logs with specific severity levels, you can use the filters to forward only the information you need. For example, a security operations analyst who investigates malware attacks might be interested only in Threat logs with the type attribute set to wildfire-virus.



## 5.1.2 Manage external services

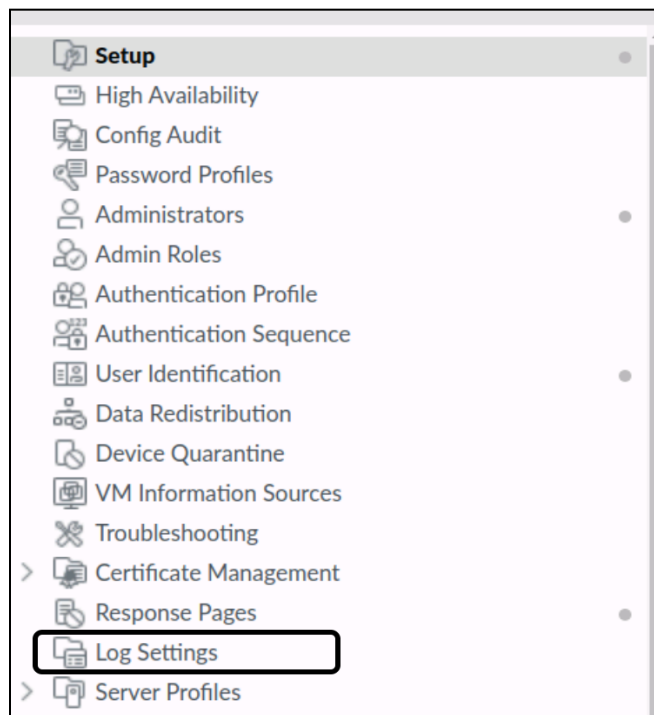
### Destination Log Types and Formatting

External forwarding supports the following types of destinations:

- SNMP traps
- Syslog
- HTTP server
- Email
- Panorama

### Filtering and Forwarding Log Events

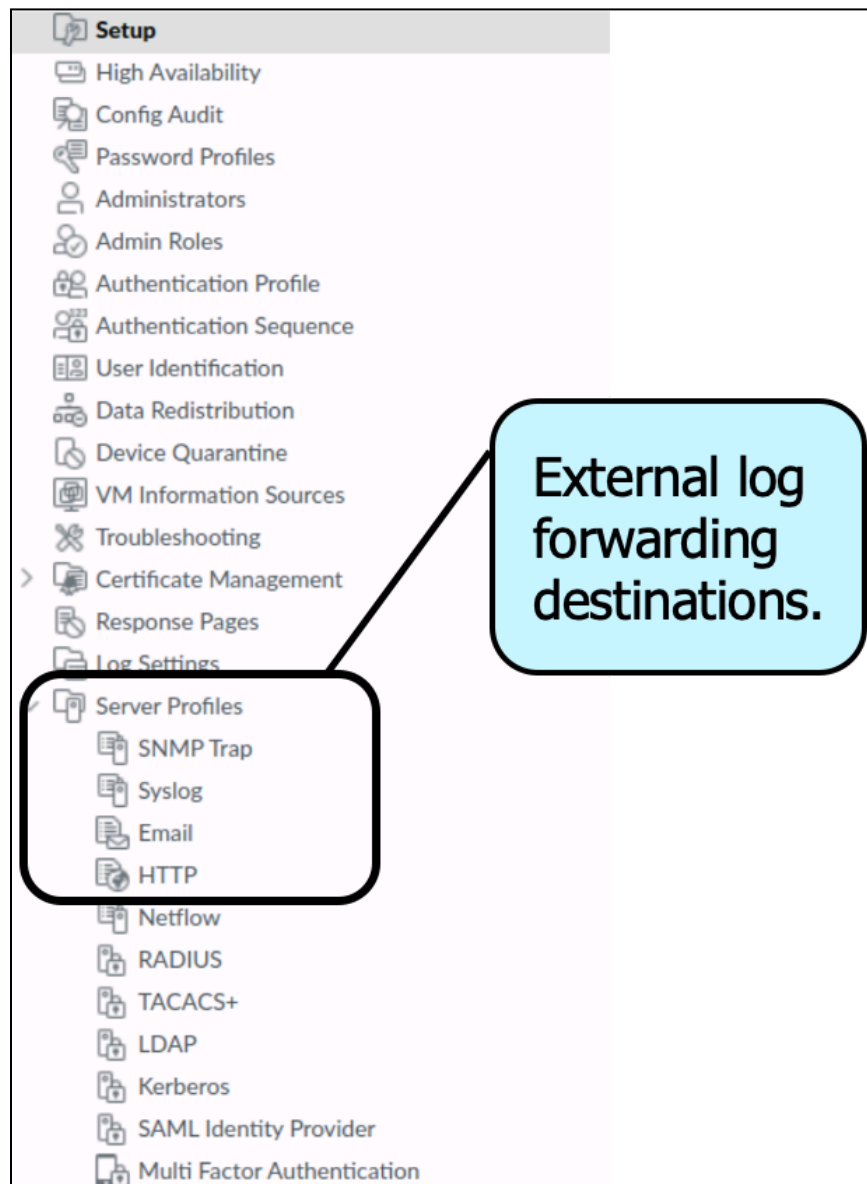
The Palo Alto Networks firewall has two primary methods to forward log events, depending on the log message type. Events associated with the examined traffic use Log Forwarding Profiles. Events that are generally related to non-traffic-specific firewall activity can be filtered and forwarded using Log Settings, found in **Device > Log Settings**.



Log forwarding of any event type can send copies of log events to external destinations supporting the following data formats:

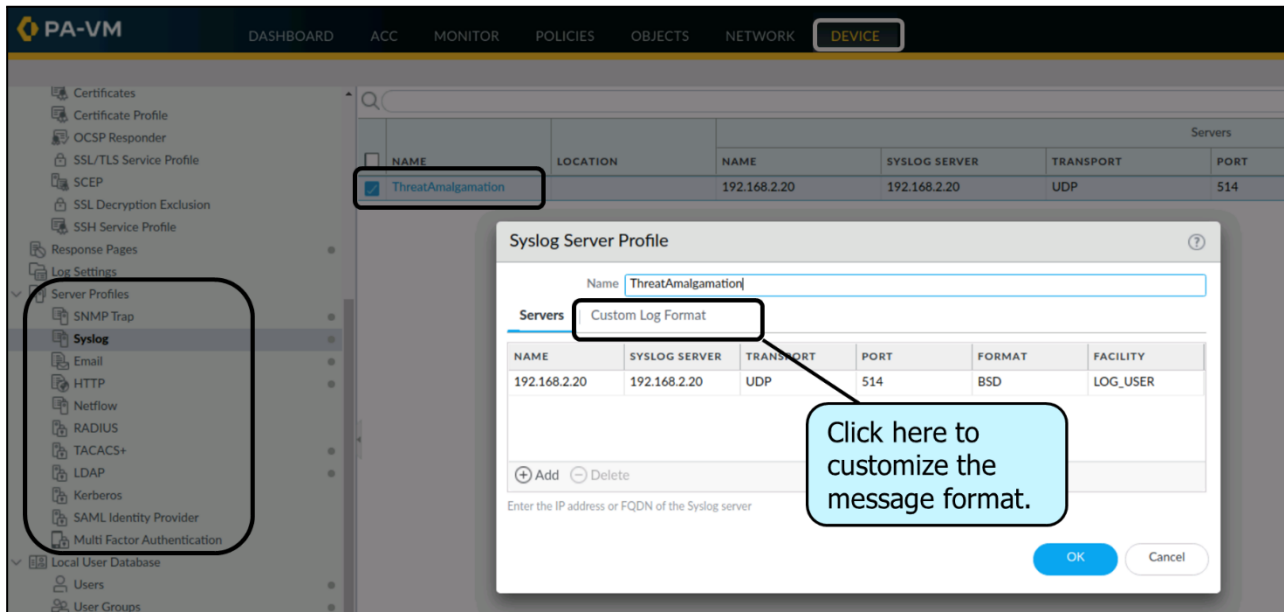
- SNMP
- Email
- Syslog
- HTTP

Each log forwarding destination is configured in the firewall with a Server Profile of the appropriate type. Navigate to **Device > Server Profiles** and create a profile for each specific destination.

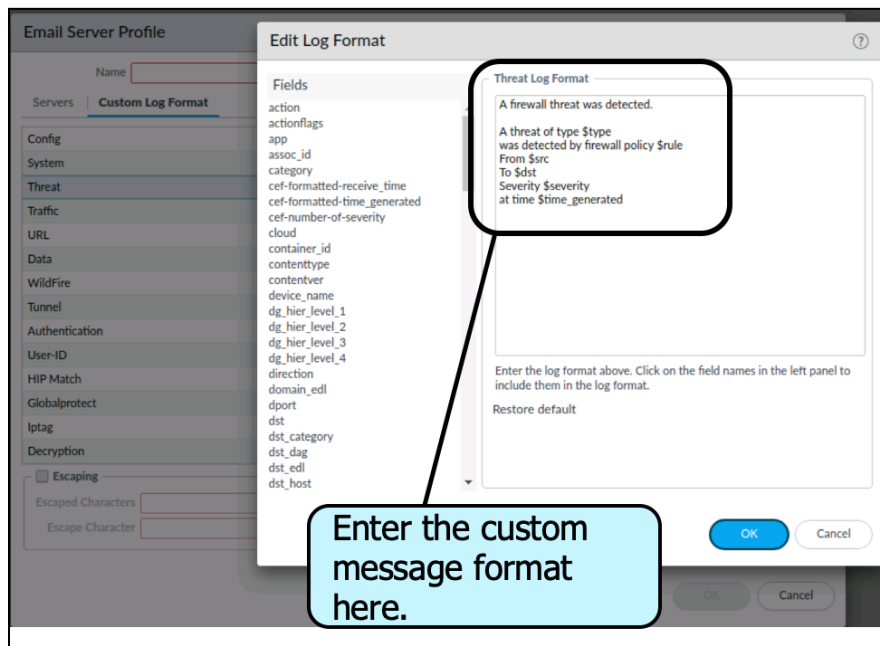


After the destination's Server Profile is created, it can be used in a Log Forwarding Profile.

All types (other than Panorama) support customization of the message format. A typical destination configuration is as follows:



Email message formats can be customized. Here is an example:

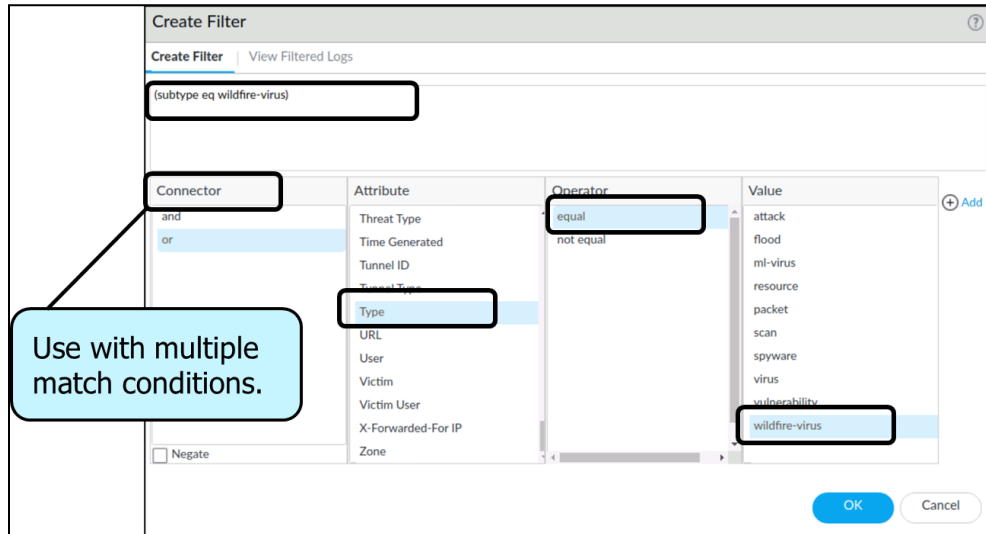


Any log event redirection causes a copy of the log event to be forwarded as specified. It is logged on the firewall as usual.

### 5.1.3 Create and manage tags

#### Automated Actions and Tagging with Log Forwarding

Log Forwarding Profiles also provide a mechanism to collect the source or destination IP address of the event and tag it. The tag can then be used to assign the address to a Dynamic Address Group that is used in a Security policy rule.



### 5.1.6 Log monitoring

A log is an automatically generated, time-stamped file that provides an audit trail for system events on the firewall or network traffic events that the firewall monitors. Log entries contain artifacts, which are properties, activities, or behaviors associated with the logged event, such as the application type or the IP address of an attacker. Each log type records information for a separate event type. For example, the firewall generates a Threat log to record traffic that matches a spyware, vulnerability, or virus signature or a DoS attack that matches the thresholds configured for a port scan or host sweep activity on the firewall.

### 5.1.4 Customize logging and reporting settings

#### Logging and Reporting Settings

Use this section to modify:

- Expiration periods and storage quotas for reports and for the following log types. The settings are synchronized across high availability (HA) pairs.
  - Logs of all the types that the firewall generates and stores locally (**Device > Setup > Management**). The settings apply to all the virtual systems on the firewall.
  - Logs that an M-Series appliance or a Panorama virtual appliance in Panorama mode generates and stores locally: System, Config, Application Statistics, and User-ID™ logs (**Panorama > Setup > Management**).
  - Logs of all the types that the Panorama virtual appliance in Legacy mode generates locally or collects from firewalls (**Panorama > Setup > Management**).
- Attributes for calculating and exporting user activity reports.
- Predefined reports created on the firewall or Panorama.

## Log Storage tab

(Panorama management server and all firewall models except PA-5200 Series and PA-7000 Series firewalls)



Panorama displays this tab if you edit the Logging and Reporting Settings (**Panorama > Setup > Management**). If you use a Panorama template to configure the settings for firewalls (**Device > Setup > Management**), see [Single Disk Storage and Multi Disk Storage tabs](#).

For each log type, specify:

- **Quota**—The **Quota**, as a percentage, allocated on the hard disk for log storage. When you change a **Quota** value, the associated disk allocation changes automatically. If the total of all the values exceeds 100%, a message appears in red and an error message will appear if you try to save the settings. If this happens, adjust the percentages so that the total is within the 100% limit.



VM-Series firewalls by default have a 0% quota allocated for **SCTP** log storage, **SCTP Summary**, **Hourly SCTP Summary**, **Daily SCTP Summary**, and **Weekly SCTP Summary**, so you must allocate some percentage for these firewalls to log SCTP information.

- **Max Days**—The length (in days) of the log expiration period (range is 1 to 2,000). The firewall or Panorama appliance automatically deletes logs that exceed the specified period. By default, there is no expiration period, which means logs never expire.

The firewall or Panorama appliance evaluates logs during creation of the logs and then deletes logs that exceed the expiration period or quota size.



Weekly summary logs can age beyond the threshold before the next deletion if they reach the expiration threshold between times when the firewall deletes logs. When a log quota reaches the maximum size, new log entries start overwriting the oldest log entries. If you reduce a log quota size, the firewall or Panorama removes the oldest logs when you commit the changes. In an HA active/passive configuration, the passive peer does not receive logs and, therefore, does not delete them unless failover occurs and the passive peer becomes active.

- **Core Files**—If your firewall experiences a system process failure, it will generate a core file that contains details about the process and why it failed. If a core file is too large for the default core file storage location (`/var/cores` partition), you can enable the `large-core` file option to allocate an alternate and larger storage location (`/opt/panlogs/cores`). A Palo Alto Networks support engineer can increase the allocated storage if needed.

To enable or disable the `large-core` file option, enter the following CLI command from configuration mode and then `commit` the configuration:

```
# set deviceconfig setting management large-core [yes|no]
```



The core file is deleted when you disable this option.

You must use SCP from operational mode to export the core file:

```
> scp export core-file large-corefile
```



Only a Palo Alto Networks support engineer can interpret the contents of the core files.

- **Restore Defaults**—Select this option to revert to the default values.

## Session Log Storage and Management Log Storage tabs

(PA-5200 Series and PA-7000 Series firewalls only)

PA-5200 Series and PA-7000 Series firewalls store management logs and session logs on separate disks. Select the tab for each set of logs and configure the settings described in [Log Storage tab](#):

- **Session Log Storage**—Select **Session Log Quota** and set the quotas and expiration periods for Traffic, Threat, URL Filtering, HIP Match, User-ID, GTP/Tunnel, SCTP, and Authentication logs, as well as Extended Threat PCAPs.
- **Management Log Storage**—Set quotas and expiration periods for System, Config, and App Stats logs, as well as for HIP Reports, Data Filtering Captures, App PCAPs, and Debug Filter PCAPs.

## Single Disk Storage and Multi Disk Storage tabs

(Panorama template only)

If you use a Panorama template to configure log quotas and expiration periods, configure the settings in one or both of the following tabs based on the firewalls assigned to the template:

- **PA-5200 Series and PA-7000 Series firewalls**—Select **Multi Disk Storage** and configure the settings in the [Session Log Storage and Management Log Storage tabs](#).



PA-5200 Series firewalls by default have a 0% quota allocated for **SCTP** log storage, **SCTP Summary**, **Hourly SCTP Summary**, **Daily SCTP Summary**, and **Weekly SCTP Summary**, so you must allocate some percentage for these firewalls to log SCTP information.

- **All other firewall models**—Select **Single Disk Storage**, select **Session Log Quota**, and configure the settings on the [Log Storage tab](#).



## Log Export and Reporting tab

Configure the following log export and reporting settings as needed:

- **Number of Versions for Config Audit**—Enter the number of configuration versions to save before discarding the oldest ones (default is 100). You can use these saved versions to audit and compare changes in configuration.
- **Number of Versions for Config Backups**—(Panorama only) Enter the number of configuration backups to save before discarding the oldest ones (default is 100).
- **Max Rows in CSV Export**—Enter the maximum number of rows that will appear in the CSV reports generated when you **Export to CSV** from the traffic logs view (range is 1 to 1,048,576; default is 65,535).
- **Max Rows in User Activity Report**—Enter the maximum number of rows that is supported for the detailed user activity reports (range is 1 to 1,048,576; default is 5,000).

## Log Export and Reporting tab (cont)

- **Average Browse Time (sec)**—Configure this variable to adjust how the browse time is calculated in seconds for the **Monitor > PDF Reports > User Activity Report** (range is 0 to 300 seconds; default is 60).

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, see [Container Pages](#).

The average browse time setting is the average time that the administrator 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.

- **Page Load Threshold (sec)**—Allows you to adjust the assumed time (in seconds) that it takes for page elements to load on the page (range is 0 to 60; default is 20). 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. The page load threshold is also used in the calculations for the **Monitor > PDF Reports > User Activity Report**.
- **Syslog HOSTNAME Format**—Select whether to use the FQDN, hostname, or IP address (IPv4 or IPv6) in the syslog message header. This header identifies the firewall or Panorama management server where the message originated.
- **Report Runtime**—Select the time of day (default is 2 a.m.) when the firewall or Panorama appliance starts generating daily scheduled reports.
- **Report Expiration Period**—Set the expiration period (in days) for reports (range is 1 to 2,000). By default, there is no expiration period, which means reports never expire. The firewall or Panorama appliance deletes expired reports nightly at 2 A.M. according to its system time.

## (Panorama only)

- **Buffered Log Forwarding from Device** (enabled by default)—Allows the firewall to buffer log entries on its hard disk (local storage) when it loses connectivity to Panorama. When the connection to Panorama is restored, the firewall forwards the log entries to Panorama; the disk space available for buffering depends on the log storage quota for the firewall model and the volume of logs that are pending roll over. If the available space is consumed, the oldest entries are deleted to allow logging of new events.



Enable **Buffered Log Forwarding from Device** to help prevent loss of logs if the connection to Panorama goes down.

- **Get Only New Logs on Convert to Primary** (disabled by default)—This option applies only to a Panorama virtual appliance in Legacy mode that writes logs to a Network File System (NFS). With NFS logging, only the primary Panorama is mounted to the NFS. Therefore, the firewalls send logs only to the active primary Panorama. This option enables you to configure firewalls to send newly generated logs only to Panorama when an HA failover occurs and the secondary Panorama resumes logging to the NFS (after it is promoted as primary). This option is typically enabled to prevent firewalls 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** (disabled by default)—This option applies only to a Panorama virtual appliance in Legacy mode. This option enables you to configure only the active Panorama to save logs to the local disk.

- **Pre-Defined Reports** (enabled by default)—Pre-defined reports for application, traffic, threat, URL Filtering, and Stream Control Transmission Protocol (SCTP) are available on the firewall and on Panorama. Pre-defined reports for SCTP are available on the firewall and Panorama after SCTP Security is enabled in **Device > Setup > Management > General Settings**.

Because the firewalls consume memory resources in generating the results hourly (and forwarding it to Panorama where it is aggregated and compiled for viewing), to reduce memory usage, you can disable the reports that are not relevant to you. To disable a report, disable this option for the report.

Click **Select All** or **Deselect All** to entirely enable or disable the generation of pre-defined reports.



Before disabling a report, verify that there isn't a Group Report or a PDF Report using it. If you disable a predefined report assigned to a set of reports, the entire set of reports will have no data.

- **Log Admin Activity** (disabled by default)—Specify whether to generate an audit log when an administrator executes an operational command in the firewall CLI or navigates through the web interface. You must first successfully configure a syslog server before you can generate and forward an audit log.
  - **Operational Commands**—Generate an audit log when an administrator executes an operational or debug command in the CLI or an operational command that is triggered from the web interface. See the [CLI Operational Command Hierarchy](#) for a full list of PAN-OS operational and debug commands.
  - **UI Actions**—Generate an audit log when an administrator navigates throughout the web interface. This includes navigation between configuration tabs, as well as between individual objects within a tab. For example, an audit log is generated when an administrator navigates from the **ACC** to the **Policies** tab. Additionally, an audit log is generated when an administrator navigates from **Objects > Addresses** to **Objects > Tags**.
  - **Syslog Server**—Select the target syslog server profile to forward audit logs.

## 5.1.5 References

- Monitoring, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring>
- Device > Setup > Management, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/device/device-setup-management>
- CLI Cheat Sheet: Networking, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-cli-quick-start/cli-cheat-sheets/cli-cheat-sheet-networking>
- Configure Log Forwarding, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/configure-log-forwarding>
- Cortex Data Lake, <https://docs.paloaltonetworks.com/cortex/cortex-data-lake>
- What Data Center Traffic to Log and Monitor, <https://docs.paloaltonetworks.com/best-practices/10-2/data-center-best-practices/data-center-best-practice-security-policy/log-and-monitor-data-center-traffic/what-data-center-traffic-to-log-and-monitor>

## 5.2 Plan and execute the process to upgrade a Palo Alto Networks system

### 5.2.1 Single firewall

The Palo Alto Networks NGFW is provided with a Single Pass Software. It processes the packet to perform actions, such as networking, User-ID, policy lookup, traffic classification with application identification (App-ID), decoding, and signature matching for detecting threats and malicious contents. Packet processing at one go or in a single pass by the Palo Alto Networks NGFW significantly reduces the overhead of packet processing.

On the contrary, other firewall vendors leverage a different type of network architecture, which produces a higher overhead when processing packets traversing the firewall. Another notable feature introduced in another Firewall vendor's NGFW is Unified Threat Management (UTM), which processes the packet and then verifies its contents. As a result, the spike in CPU overhead affects the latency and throughput of the Firewalls, and leads to a degradation in performance.

The Single Pass software is designed to achieve two key parameters.

- First, the Single Pass Software performs operation per packet. When a packet is processed using this mechanism, functions such as policy lookup, application identification and decoding, and signature matching for all the threats and content are performed just once.
- Second, the packet processed in the Single Pass software is stream-based and uses uniform signature matching to detect and block threats. Single Pass does not use separate engines and signature sets and file proxies required for file download prior to scanning; the single pass software in our next generation firewalls scans packets once in stream-based manner to avoid latency and throughput.

This content processing by the Single Pass software enables high throughput and low latency with all of the security functions active. The software also offers the additional feature of a single fully-integrated policy, enabling easier management of enterprise network security.

### 5.2.2 High availability pairs

You can configure two Palo Alto Networks firewalls as an HA pair or up to 16 firewalls as peer members of an HA cluster. The peers in the cluster can be HA pairs or standalone firewalls. HA allows you to minimize downtime by making sure that an alternate firewall is available in case a peer firewall fails. The firewalls in an HA pair or cluster use dedicated or in-band HA ports on the firewall to synchronize data — network, object, and policy configurations — and maintain state information. Firewall-specific configuration, such as management interface IP address or administrator profiles, HA-specific configuration, log data, and Application Command Center (ACC) information, is not shared between peers.



For a consolidated application and log view across an HA pair, you must use Panorama, the Palo Alto Networks centralized management system. When a failure occurs on a firewall in an HA pair or cluster and a peer firewall takes over the task of securing traffic, the event is called a “failover.” The conditions that trigger a failover are:

- One or more of the monitored interfaces fail.
- One or more of the destinations specified on the firewall cannot be reached.
- The firewall does not respond to heartbeat polls.
- A critical chip or software component fails, known as packet path health monitoring.

The Palo Alto Networks firewalls support stateful active/passive or active/active HA with session and configuration synchronization, with a few exceptions:

- The VM-Series firewall on Azure and VM-Series firewall on AWS only support active/passive HA. When you deploy the firewall with the Amazon Elastic Load Balancing (ELB) service on AWS, it does not support HA (in this case, ELB service provides the failover capabilities).
- The VM-Series firewall on Google Cloud Platform does not support traditional HA.

### 5.2.3 Panorama push

#### **Upgrading Firewalls Under Panorama Management**

Firewalls managed by Panorama can get dynamic updates from Panorama, including scheduled updates. PAN-OS upgrades can also be managed from Panorama. A pair of Panorama instances can be used to download software updates. One Panorama with a trusted internet connection can transfer updates to an SCP server while the second Panorama deployed in an isolated network can use the SCP server as a software update server. The second Panorama can then download any updates and then send them to all the managed devices.

#### **HA Cluster Firewall Updates Managed by Panorama**

Panorama treats the managed firewalls in HA pairs as individual firewalls for software update purposes.

### 5.2.4 Dynamic updates

Palo Alto Networks frequently publishes the updates that the firewall can use to enforce the security policy, without requiring to upgrade the PAN-OS software or change the firewall configuration. These updates equip the firewall with the very latest security features and threat intelligence.

Except for application updates and some antivirus updates—which any firewall can receive—dynamic content updates available to you might depend on your subscriptions. You can set a schedule for each dynamic content update to define the frequency at which the firewall checks for and downloads or installs new updates.

## 5.2.5 References

- Palo Alto Firewall Architecture, <https://networkinterview.com/palo-alto-firewall-architecture/#:~:text=Palo%20Alto%20Firewal%20Architecture%20is%20based%20upon%20an.network%20security%20integrated%20with%20remarkably%20features%20and%20technology>
- HA Overview, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/ha-overview>
- Upgrade Firewalls Using Panorama, <https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-upgrade/upgrade-pan-os/upgrade-the-firewall-pan-os/upgrade-firewalls-using-panorama>
- Automatic Content Updates Through Offline Panorama, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-upgrade/upgrade-panorama/install-content-and-software-updates-for-panorama/install-updates-automatically-for-panorama-without-an-internet-connection>
- Dynamic Content Updates, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-upgrade/software-and-content-updates/dynamic-content-updates>

## 5.3 Manage HA functions

### 5.3.1 Link monitoring

#### Settings Related to Critical HA Functions

An HA pair configuration is created when two firewalls are placed in a group and have their configurations synchronized to prevent a single point of failure on the network. A heartbeat connection between the firewall peers ensures seamless failover if a peer becomes non-operational. Configure two firewalls in an HA pair to provide redundancy and ensure business continuity.

An HA cluster can be configured with up to 16 firewalls or HA pairs acting in an all-active manner. HA clusters require an HA4 link to synchronize session state information, and they use link and path monitoring to determine the up/down state of cluster members. HA4 and potential HA4 backup links determine HA cluster member functionality.

#### HA Functionality

Network monitoring applications use SNMP to query network components such as the NGFW. The firewall has additional information that is specific to HA. You can monitor the dedicated HA1, HA2, HA2 backup, and HA3 interfaces. Use the IF-MIB and the interface's Management Information Base (MIB) to see the SNMP statistics for dedicated HA2 interfaces.

Panorama includes Managed Device Health Monitoring, which displays limited HA status information in the summary display in the Panorama management web interface.

### 5.3.2 Path monitoring

You can specify a destination IP group of IP addresses for the firewall to monitor. The firewall monitors the full path through the network to mission-critical IP addresses by using ICMP pings to verify IP address reachability. The default interval for pings is 200ms. An IP address is considered unreachable when 10 consecutive pings (the default value) fail. You specify the failure condition for the IP addresses in a destination IP group—Any IP address unreachable or All IP addresses unreachable in the group. You specify multiple destination IP groups for a path group for a virtual wire, VLAN, or virtual router. You specify the failure condition of destination IP groups in a path group—Any or All, which constitutes a path group failure. You can configure multiple virtual wire path groups, VLAN path groups, and virtual router path groups.

You also determine the global failure condition—Any path group fails or All path groups fail, which determines when a failover is triggered. The default behavior is that Any of the IP addresses becoming unreachable in Any destination IP group in Any virtual wire, VLAN, or virtual router path group causes the firewall to change the HA state to non-functional (or to tentative state in active/active mode) to indicate the failure of a monitored object.

### 5.3.3 HA links

The firewalls in an HA pair use HA links to synchronize data and maintain state information. Some models of the firewall have dedicated HA ports—Control link (HA1) and Data link (HA2), while others require you to use the in-band ports as HA links.

- For firewalls with dedicated HA ports, use these ports to manage communication and synchronization between the firewalls.
- For firewalls without dedicated HA ports such as the PA-220 and PA-220R firewalls, as a best practice use the management port for the HA1 port, and use the dataplane port for the HA1 backup.
- For firewalls without dedicated HA ports, decide which ports to use for HA1 and HA1 backup based on your environment and understanding which are the least used and least congested. Assign HA1 to the best interface and HA1 backup to the other one.

HA peers in an HA cluster can be a combination of standalone members and HA pairs. HA cluster members use an HA4 link and HA4 backup link to perform session state synchronization. HA1 (control link), HA2 (data link), and HA3 (packet-forwarding link) are not supported between cluster members that aren't HA pairs.

## HA LINKS AND BACKUP LINKS

### DESCRIPTION

#### Control Link

The HA1 link is used to exchange hellos, heartbeats, and HA state information, and management plane sync for routing, and User-ID information. The firewalls also use this link to synchronize configuration changes with its peer. The HA1 link is a Layer 3 link and requires an IP address.

ICMP is used to exchange heartbeats between HA peers.

Ports used for HA1—TCP port 28769 and 28260 for clear text communication; port 28 for encrypted communication (SSH over TCP).

If you enable encryption on the HA1 link, you can also refresh HA1 SSH keys and configure key options.

#### Data Link

The HA2 link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between firewalls in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active or active-primary firewall to the passive or active-secondary firewall. The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.

Ports used for HA2—The HA data link can be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets.

## HA1 and HA2 Backup Links

Provide redundancy for the HA1 and the HA2 links. In-band ports can be used for backup links for both HA1 and HA2 connections when dedicated backup links are not available. Consider the following guidelines when configuring backup HA links:

- The IP addresses of the primary and backup HA links must not overlap each other.
- HA backup links must be on a different subnet from the primary HA links.
- HA1-backup and HA2-backup ports must be configured on separate physical ports. The HA1-backup link uses port 28770 and 28260.
- PA-3200 Series firewalls don't support an IPv6 address for the HA1-backup link; use an IPv4 address.

Palo Alto Networks recommends enabling heartbeat backup (uses port 28771 on the MGT interface) if you use an in-band port for the HA1 or the HA1 backup links.

## Packet-Forwarding Link

In addition to HA1 and HA2 links, an active/active deployment also requires a dedicated HA3 link. The firewalls use this link for forwarding packets to the peer during session setup and asymmetric traffic flow. The HA3 link is a Layer 2 link that uses MAC-in-MAC encapsulation. It does not support Layer 3 addressing or encryption. PA-7000 Series firewalls synchronize sessions across the NPCs one-for-one. On PA-800 Series, PA-3200 Series, PA-3400 Series, PA-5200 Series, and PA-5400 Series firewalls, you can configure aggregate interfaces as an HA3 link. The aggregate interfaces can also provide redundancy for the HA3 link; you cannot configure backup links for the HA3 link. On PA-3200 Series, PA-3400 Series, PA-5200 Series, PA-5400 Series, and PA-7000 Series firewalls, the dedicated HSCI ports support the HA3 link. The firewall adds a proprietary packet header to packets traversing the HA3 link, so the MTU over this link must be greater than the maximum packet length forwarded.

HA4 Link and HA4  
Backup Link

The HA4 link and HA4 backup link perform session cache synchronization among all HA cluster members having the same cluster ID. The HA4 link between cluster members detects connectivity failures between cluster members by sending and receiving Layer 2 keepalive messages. View the status of the HA4 and HA4 backup links on the firewall dashboard.

### 5.3.4 Failover

When a failure occurs on one firewall and the peer in the HA pair (or a peer in the HA cluster) takes over the task of securing traffic, the event is called a failover. A failover is triggered, for example, when a monitored metric on a firewall in the HA pair fails. The metrics that the firewall monitors for detecting a firewall failure are:

- **Heartbeat Polling and Hello messages**
- **Link Monitoring**
- **Path Monitoring**

A failover also occurs when the administrator suspends the firewall or when preemption occurs.

On PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls, a failover can occur when an internal health check fails. This health check is not configurable and is enabled to monitor the critical components, such as the FPGA and CPUs. Additionally, general health checks occur on any platform, causing failover.

The following describes what occurs in the event of a failure of a Network Processing Card (NPC) on a PA-7000 Series firewall that is a member of an HA cluster:

- If the NPC that is being used to hold the HA clustering session cache (a copy of the other members' sessions) goes down, the firewall goes non-functional. When this occurs, the session distribution device (such as a load balancer) must detect that the firewall is down and distribute session load to the other members of the cluster.
- If the NPC of a cluster member goes down and no link monitoring or path monitoring was enabled on that NPC, the PA-7000 Series firewall member will stay up, but with a lower capacity because one NPC is down.
- If the NPC of a cluster member goes down and link monitoring or path monitoring was enabled on that NPC, the PA-7000 Series firewall will go non-functional and the session distribution device (such as a load balancer) must detect that the firewall is down and distribute session load to the other members of the cluster.

### 5.3.5 Active/active and active/passive

#### HA Pair Modes

The Palo Alto Networks firewalls support stateful active/passive or active/active HA with session and configuration synchronization, with a few exceptions:

- The VM-Series firewall in Amazon Web Services supports active/passive HA only; if the firewall is deployed with Amazon Elastic Load Balancing (ELB), it does not support HA. (In this case, ELB provides the failover capabilities.)
- The VM-Series firewall in Microsoft Azure supports active/passive HA in PAN-OS 9x or later.
- The VM-Series firewall on Google Cloud Platform does not support traditional HA.

Public cloud deployments of the VM-Series firewalls also are supported in each vendor's version of a "scaled" implementation, thus allowing virtual firewalls to share the traffic load through a deployment of parallel firewall instances and providing the option to create or remove firewall instances with changing traffic loads. These deployments all include the cloud vendor's load balancer deployed in front of the firewall "scale set" to manage the spreading of the traffic across the available firewalls. This same deployment practice also creates an HA scenario in the sense that failing firewall instances can be removed from the scale set automatically using various detection abilities within the load balancer. A limitation of the scale set methods of HA is that there is typically no synchronization between firewalls and so failovers are disruptive because the existing sessions are terminated.

#### Active/Passive Pairs

Active/passive HA usually is the recommended deployment method. One firewall actively manages traffic while the other is synchronized and ready to transition to the active state if a failure occurs. In this mode, both firewalls share the same configuration settings and one actively manages traffic until a path, link, or system failure occurs. When the active firewall fails, the passive firewall transitions to the active state, takes over seamlessly, and enforces the same policies to maintain network security. The firewalls synchronize their session state tables, thus allowing the passive partner to become active and continue servicing active sessions at failover. Active/passive HA is supported in the virtual wire, Layer 2, and Layer 3 deployments.

Active/passive usually is much easier to manage because one firewall is handling traffic and both firewalls share the same traffic interface configuration.

#### Active/Active Pairs

In active/active HA, both firewalls in the pair are active and processing traffic. They work synchronously to handle session setup and session ownership. Both firewalls individually maintain session tables and routing tables and synchronize with each other. Active/active HA is supported in both virtual wire and Layer 3 deployments. In active/active HA mode, the firewall HA interfaces cannot receive addresses via DHCP. Furthermore, only the active-primary firewall's traffic interface can function as a DHCP relay. The active-secondary firewall that receives the DHCP broadcast packets drops them.


In a Layer 3 deployment of HA active/active mode, you can assign floating IP addresses that move from one HA firewall to the other if a link or firewall fails. The interface on the firewall that owns the floating IP address responds to ARP requests with a virtual MAC address.

Floating IP addresses are recommended when you need functionality such as the Virtual Router Redundancy Protocol. Floating IP addresses also can be used to implement VPNs and the source NAT, thus allowing for persistent connections if a firewall offering those services fails.

Each HA firewall interface has its own IP address and floating IP address. The interface IP address remains local to the firewall, but the floating IP address moves between firewalls upon firewall failure. You configure the end hosts to use a floating IP address as its default gateway, thus allowing load-balance of traffic to the two HA peers. You can also use external load balancers to load-balance traffic.

If a link or firewall fails or a path monitoring event causes a failover, the floating IP address and virtual MAC address move over to the functional firewall. (In the figure that follows, each firewall has two floating IP addresses and virtual MAC addresses; they all move over if the firewall fails.) The functioning firewall sends a gratuitous ARP to update the MAC tables of the connected switches to inform them of the change in the floating IP address and MAC address ownership for redirecting traffic to itself.

After the failed firewall recovers, by default the floating IP address and virtual MAC address move back to the firewall with the Device ID (0 or 1) to which the floating IP address is bound. More specifically, after the failed firewall recovers, it becomes online. The currently active firewall determines that the firewall is back online and checks whether the floating IP address that it is handling belongs natively to itself or to the other firewall. If the floating IP address was originally bound to the other Device ID, the firewall automatically gives it back. An example of a floating IP deployment follows:



Each firewall in the HA pair creates a virtual MAC address for each of its interfaces that has a floating IP address or ARP load-sharing IP address.



### 5.3.6 HA interfaces

#### HA Links and Backup Links

The firewalls in an HA pair and cluster use HA links to synchronize data and maintain state information. Some firewall models have dedicated HA ports—control link (HA1) and data link (HA2)—while others require you to use the in-band ports as HA links. Firewalls in an HA cluster use an in-band Layer 3 HA4 interface for cluster session synchronization as follows:

- For firewalls with dedicated HA ports, use these ports to manage communication and synchronization between the firewalls.
- For firewalls without dedicated HA ports, use a data plane port for the HA port and use the management port as the HA1 backup.

A recommended best practice is the implementation of backup HA paths because the HA ports synchronize data that is critical to proper HA failover. In-band ports can be used as backup links for the HA1, HA2, and HA3 connections when the dedicated backup links are not available. Consider the following guidelines when you configure backup HA links:

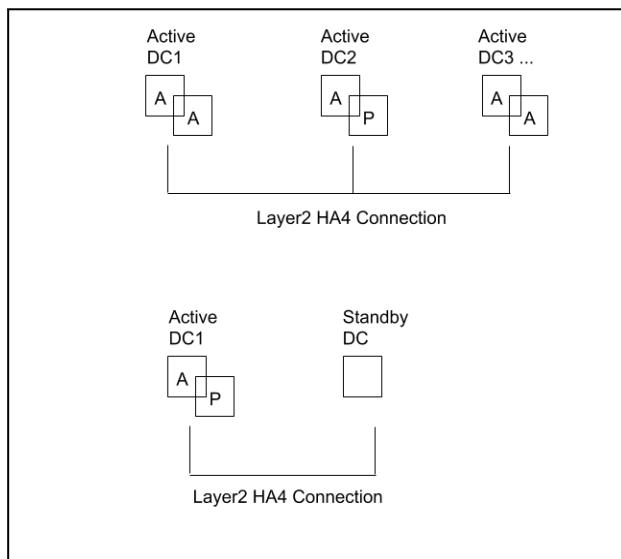
- The IP addresses of the primary and backup HA links must not overlap.
- The HA backup links must be on a subnet that is different from the primary HA links.
- The HA1-backup and HA2-backup ports must be configured on separate physical ports. The HA1-backup link uses ports 28770 and 28260.

### 5.3.7 Clustering

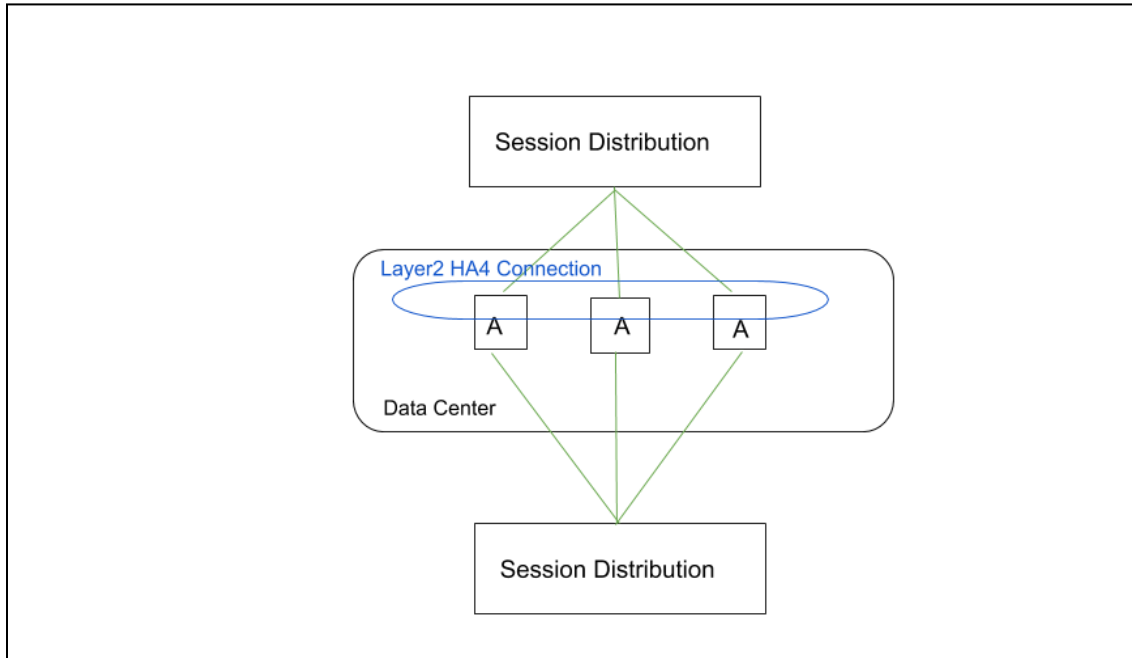
Many Palo Alto networks firewall models support session state synchronization among firewalls in an HA cluster of up to 16 firewalls. The HA cluster peers synchronize sessions to protect against data center or large security inspection point failure with horizontally scaled firewalls. In the case of a network outage or a down firewall, the sessions fail over to a different firewall in the cluster.

Such synchronization is especially helpful in the following use cases:

- When HA peers are spread across multiple data centers
- When one data center is active and the other is on standby



A third HA clustering use case is horizontal scaling, in which you add HA cluster members to a single data center to scale security and ensure session survivability.



HA clusters support a Layer 3 or virtual wire deployment. HA peers in the cluster can be a combination of HA pairs and standalone cluster members. In an HA cluster, all the members are considered active; there is no concept of passive firewalls except for HA pairs, which can keep their active/passive relationship after you add them to an HA cluster.

All cluster members share session state. When a new firewall joins an HA cluster, that triggers all of the firewalls in the cluster to synchronize all the existing sessions. The HA4 and HA4 backup connections are the dedicated cluster links that synchronize session state among all the cluster members with the same cluster ID. The HA4 link between cluster members detects connectivity failures between cluster members. HA1 (control link), HA2 (data link), and HA3 (packet-forwarding link) are not supported between cluster members that aren't HA pairs.

For a normal session that has not failed over, only the firewall that is the session owner creates a traffic log. For a session that failed over, the new session owner (the firewall that receives the failed over traffic) creates the traffic log.

The firewall models that support HA clustering and the maximum number of members supported per cluster are described in the following table:

FIREWALL MODEL	NUMBER OF MEMBERS SUPPORTED PER CLUSTER
PA-3200 Series	6
PA-5200 Series	16
PA-7000 Series firewalls that have at least one of the following cards: PA-7000-100G-NPC, PA-7000-20GQXM-NPC, and PA-7000-20GXM-NPC	PA-7080: 4 PA-7050: 6
VM-300	6
VM-500	6
VM-700	16

### 5.3.8 Election setting

Specify or enable the following settings:

- **Device Priority:** Uses a priority value to identify the active firewall. The firewall with the lower value (higher priority) becomes the active firewall (range is 0 to 255) when the preemptive capability is enabled on both firewalls in the pair.
- **Preemptive:** Enables the higher-priority firewall to resume active (active/passive) or active-primary (active/active) operation after recovering from a failure. You must enable the preemption option on both firewalls for the higher-priority firewall to resume active or active-primary operation upon recovery after a failure. If this setting is disabled, then the lower-priority firewall remains active or active-primary even after the higher-priority firewall recovers from a failure.
- **Heartbeat Backup:** Uses the management ports on the HA firewalls to provide a backup path for heartbeat and Hello messages. The management port IP address will be shared with the HA peer through the HA1 control link. No additional configuration is required.

### 5.3.9 References

- HA Concepts, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/ha-concepts>
- Failover, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/ha-concepts/failover>
- What is HA-Lite on Palo Alto Networks PA-200?, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIUzCAK>
- HA Clustering Overview, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/ha-clustering-overview>
- HA Links and Backup Links, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/ha-concepts/ha-links-and-backup-links>
- Set Up Active/Passive HA, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/set-up-active-passive-ha>
- Configure HA Clustering, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/configure-ha-clustering>
- HA Clustering Best Practices and Provisioning, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/ha-clustering-best-practices-and-provisioning>
- SNMP Support, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/snmp-monitoring-and-traps/snmp-support>
- Monitor Statistics Using SNMP, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/snmp-monitoring-and-traps/monitor-statistics-using-snmp>
- Supported MIBs, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/snmp-monitoring-and-traps/supported-mibs>
- Monitor Device Health, <https://docs.paloaltonetworks.com/panorama/11-0/panorama-admin/manage-firewalls/device-monitoring-on-panorama/monitor-device-health>
- Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/set-up-active-active-ha/determine-your-activeactive-use-case/use-case-configure-activeactive-ha-with-floating-ip-address-bound-to-active-primary-firewall>
- Information Synchronized in an HA Pair, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIXGCA0>
- What Settings Don't Sync in Active/Passive HA?, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/high-availability/reference-ha-synchronization/what-settings-dont-sync-in-activepassive-ha>
- HA General Settings, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/device/device-high-availability/ha-general-settings>

## Domain 6: Troubleshooting

### 6.1 Troubleshoot site-to-site tunnels

Palo Alto Networks firewall troubleshooting involves a wide range of specific knowledge that depends on the type of issue involved. This section introduces a few principal tools and methods available for troubleshooting. The end of this section includes references for other tools and topics. Dedicated training classes for firewall troubleshooting are also available from Palo Alto Networks and Training Partners.

#### 6.1.1 IPsec

The IPsec tunnel configuration allows you to authenticate and/or encrypt the data (IP packet) as it traverses the tunnel.

If you are setting up the firewall to work with a peer that supports policy-based VPN, you must define Proxy IDs. Devices that support policy-based VPN use specific security rules/policies or access-lists (source addresses, destination addresses and ports) for permitting interesting traffic through an IPsec tunnel. These rules are referenced during quick mode/IKE phase 2 negotiation and are exchanged as Proxy-IDs in the first or the second message of the process. So, if you are configuring the firewall to work with a policy-based VPN peer, then for a successful phase 2 negotiation, you must define the Proxy-ID so that the setting on both peers is identical. If the Proxy-ID is not configured because the firewall supports route-based VPN, the default values used as Proxy-ID are source ip: 0.0.0.0/0, destination ip: 0.0.0.0/0, and application: any; and when these values are exchanged with the peer, it results in a failure to set up the VPN connection.

#### 6.1.2 GRE

Considerations for troubleshooting GRE tunnels:

- Check GRE Tunnel Status:
  - From CLI run command shown below
  - Verify "tunnel interface state" field. The value should be "Up"
  - Use the "show counter global" command with filter flow\_gre (show counter global filter value all | match flow\_gre)
  - The counter should display value of decap\_success and encap\_success increase when traffic is being sent through the tunnel
- Check the System log:
  - Run "show log system" with the filter of tunnel name. The tunnel should be up with no flaps.
- Check the GRE session
  - Filter by using session filter protocol 47
  - Find the session ID and filter for session ID
- Verify "session terminate tunnel" value which must be "True" and verify ingress and egress interface

### 6.1.3 One-to-one and one-to-many tunnels

The following are recommendations for troubleshooting IPSec connections:

#### Phase 1 issues:

- To rule out ISP-related issues, try pinging the peer IP from the PA external interface. Ensure that pings are enabled on the peer's external interface.
- If pings have been blocked per security requirements, see if the other peer is responding to the main/aggressive mode messages, or the DPDs. Check for the responses of the "Are you there?" messages from the peer in the system logs under the Monitor tab or under ikemgr logs.
- Check that the IKE identity is configured correctly.
- Check that the policy is in place to permit IKE and IPSec applications. Usually this policy is not required if there is no clean-up rule configured on the box. If a clean-up rule is configured, the policy is usually configured from the external zone to the external zone.
- Check that proposals are correct. If incorrect, logs about the mismatch can be found under the system logs, or by using the following CLI command:  

```
> less mp-log ikemgr.log
```
- Check that preshared key is correct. If incorrect, logs about the mismatch can be found under the system logs, or by using the following CLI command:  

```
> less mp-log ikemgr.log
```
- Take packet captures to analyze the traffic. Use filters to narrow the scope of the captured traffic.
- Useful CLI commands:  

```
> show vpn ike-sa gateway <name>
> test vpn ike-sa gateway <name>
> debug ike stat
```

#### Advanced CLI commands:

- For detailed logging, turn on the logging level to debug:  

```
> debug ike global on debug
> less mp-log ikemgr.log
```
- To view the main/aggressive and quick mode negotiations, it is possible to turn on pcaps for capturing these negotiations. Messages 5 and 6 onwards in the main mode and all the packets in the quick mode have their data payload encrypted:  

```
> debug ike pcap on
> view-pcap no-dns-lookup yes no-port-lookup yes debug-pcap ikemgr.pcap
```
- Turn off debugs  

```
> debug ike pcap off
```
- Configuring packet filter and captures restricts pcaps only to the one worked on, debug IKE pcap on shows pcaps for all VPN traffic.
- To check if NAT-T is enabled, packets will be on port 4500 instead of 500 from the 5th and 6th messages of main mode.
- Check if vendor id of the peer is supported on the Palo Alto Networks device and vice-versa.

## Phase 2 issues:

- Check if the firewalls are negotiating the tunnels, and ensure that 2 unidirectional SPIs exist:

```
> show vpn ipsec-sa
  > show vpn ipsec-sa tunnel <tunnel.name>
```

- Check if proposals are correct. If incorrect, logs about the mismatch can be found under the system logs under the monitor tab, or by using the following command:

```
> less mp-log ikemgr.log
```

- Check if PFS is enabled on both ends. If incorrect, logs about the mismatch can be found under the system logs under the monitor tab, or by using the command:

```
> less mp-log ikemgr.log
```

- Check the proxy-id configuration. This is usually not required when the tunnel is between two Palo Alto Networks firewalls, but when the peer is from another vendor, IDs usually need to be configured. A mismatch would be indicated under the system logs, or by using the command:

```
> less mp-log ikemgr.log
```

- Useful CLI commands:

```
> show vpn flow name <tunnel.id/tunnel.name>
  > show vpn flow name <tunnel.id/tunnel.name> | match bytes
```

- Check if encapsulation and decapsulation bytes are increasing. If the firewall is passing traffic, then both values should be increasing.

```
> show vpn flow name <tunnel.id/tunnel.name> | match bytes
```

- If encapsulation bytes are increasing and decapsulation is constant, then the firewall is sending but not receiving packets.
- Check to see if a policy is dropping the traffic, or if a port translating device in front of PAN might be dropping the ESP packets.

```
> show vpn flow name <tunnel.id/tunnel.name> | match bytes
```

- If decapsulation bytes are increasing and encapsulation is constant, then the firewall is receiving but not transmitting packets.
- Check to see if a policy is dropping the traffic:

```
> test routing fib-lookup virtual-router default ip <destination IP>
```

```
-----
runtime route lookup
-----
```

```
virtual-router: default
destination:    10.5.1.1
               result:    interface tunnel.1
```

```
> show routing route
```

```
> test vpn ipsec-sa tunnel <name>
```

- **Advanced CLI Commands:**

```
> debug ike global on debug
> less mp-log ikemgr.log
> debug ike pcap on
> view-pcap no-dns-lookup yes no-port-lookup yes debug-pcap ikemgr.pcap
> debug ike pcap off
```

If tunnels are up but traffic is not passing through the tunnel:

- Check security policy and routing.
- Check for any devices upstream that perform port-and-address-translations. Since ESP is a layer 3 protocol, ESP packets do not have port numbers. When such devices receive ESP packets, there is a high possibility they may silently drop them, because they do not see the port numbers to translate.
- If necessary, apply debug packet filters, captures, or logs to isolate the issue where the traffic is getting dropped.

#### 6.1.4 Route-based versus policy-based remote hosts

##### **Policy-based VPNs**

- The IPsec tunnel is invoked during policy lookup for traffic matching the interesting traffic.
- There are no tunnel interfaces. The remote end of the interesting traffic has a route pointed out through the default gateway.
- As there are no tunnel interfaces, we cannot route traffic over the VPNs.
- The policies and access lists configured for the interesting traffic serve as the proxy-IDs for the tunnels.
- Firewalls that support policy-based VPNs include Juniper SRX, Juniper NetScreen, Cisco ASA, and Check Point.

##### **Route-based VPNs**

- The IPsec tunnel is invoked during route lookup for the remote end of the proxy-IDs.
- The remote end of the interesting traffic has a route pointing through the tunnel interface.
- These do support routing over the VPNs.
- Proxy-IDs are configured as part of the VPN setup.
- Firewalls that support route-based VPNs include Palo Alto Networks firewalls, Juniper SRX, Juniper NetScreen, and Check Point.

The Palo Alto Networks firewalls do not support policy-based VPNs. Policy-based VPNs have specific Security rules, policies, or access lists (source addresses, destination addresses, and ports) configured to permit interesting traffic through the IPsec tunnels. These rules are referenced during the quick mode/IPsec phase 2 and are exchanged in the first or second messages as the proxy-IDs. If the Palo Alto Networks firewall is not configured with the proxy-ID settings, the ikemgr daemon sets the proxy-ID with the default values of source ip 0.0.0.0/0, destination ip 0.0.0.0/0, and application any. These values are exchanged with the peer during the first or second message of the quick mode. A successful phase 2 negotiation requires not only that the security proposals match, but also that the proxy-IDs on either peer be a mirror image of each other.



It is mandatory to configure proxy-IDs whenever you establish a tunnel between a Palo Alto Networks firewall and the firewalls configured for policy-based VPNs.

### 6.1.5 Tunnel monitoring

Tunnel monitoring is a feature that verifies whether traffic is successfully passing across the IPsec tunnel by sending a ping down the tunnel to the configured destination.

For a VPN tunnel, you can check the connectivity to a destination IP address across the tunnel. The network monitoring profile on the firewall allows you to verify connectivity (using ICMP) to a destination IP address or a next hop at a specified polling interval and to specify an action on failure to access the monitored IP address.

If the destination IP is unreachable, you should either configure the firewall to wait for the tunnel to recover or configure automatic failover to another tunnel. In either case, the firewall generates a system log that alerts you to a tunnel failure and renegotiates the IPsec keys to accelerate recovery.

### 6.1.6 References

- VPN Deployments, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/vpn-deployments>
- Site-to-Site VPN Overview, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/site-to-site-vpn-overview>
- Large Scale VPN (LSVPN), <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/large-scale-vpn-lsvpn>
- How to Troubleshoot IPsec VPN Connectivity Issues, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000ClivCAC>
- Tunnel Monitoring, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/vpns/site-to-site-vpn-concepts/tunnel-monitoring>

## 6.2 Troubleshoot interfaces

### 6.2.1 Transceivers

You can monitor the status of transceivers in the physical appliance or device to enable easier installation and troubleshooting. Diagnostics that can be viewed include transmitted bias current, transmitted power, received power, transceiver temperature, and power supply voltage. The following devices support transceiver monitoring:

- PA-800 Series
- PA-3200 Series
- PA-5200 Series
- PA-5450 Firewall
- PA-7000 Series

Use the CLI to run transceiver monitoring.

## 6.2.2 Settings

### Troubleshooting and Configuring Interface Components

PAN-OS supports various interface configuration options. There are two general types of network interfaces on a firewall: traffic ports and the management port.

#### Traffic Ports

Traffic ports provide multiple configuration options and the ability to pass traffic through to other ports via traffic-handling objects, such as virtual routers, virtual wires, and VLANs.

#### Management Port

The management port is isolated from internal connectivity for security purposes. If the management port requires internet access, its traffic must be routed out of the firewall and through other network infrastructure that provides this connectivity. The traffic is often routed back to a traffic port on the firewall, requiring appropriate Security policy rules for access. This traffic is treated like any other transit traffic, which means that you must configure Security policy rules to allow the traffic to pass.

#### Troubleshooting Tools

There are several tools for troubleshooting traffic flow through the firewall. A best practice in troubleshooting is to separate general connectivity issues from security issues. Connectivity issues should be resolved before security processing is evaluated.

The web interface provides several tools. The path **Monitor > Logs > Traffic** provides session summary information. Log entries for traffic are generated as specified in the Security policy rules. The typical configuration specifies that log entries are created when a session ends. Use the magnifying glass icon to examine this log entry for detail:

The screenshot shows the PA-VM web interface. The top navigation bar includes Dashboard, ACC, MONITOR, POLICIES, OBJECTS, NETWORK, and DEVICE. The left sidebar lists various logs and tools. The main area displays a log table with columns for RECEIVE TIME, TYPE, ACTION, RULE, RULE UUID, BY, SEVERITY, CATEG, URL CATEG, LIST, VERD, URL, and FILE NAME. A 'Detailed Log View' pop-up window is open, showing session details for a log entry. A magnifying glass icon is highlighted over a log entry in the table, and a callout box points to it with the text 'Click for detailed information for this log entry.'

RECEIVE TIME	TYPE	ACTION	RULE	RULE UUID	BY...	SEVERI...	CATEG...	URL CATEG...	LIST	VERDL...	URL	FILE NAME
2020/05/18 16:36:26	end	allow	Users_L...	7a2d1...	484		any					

**Detailed Log View**

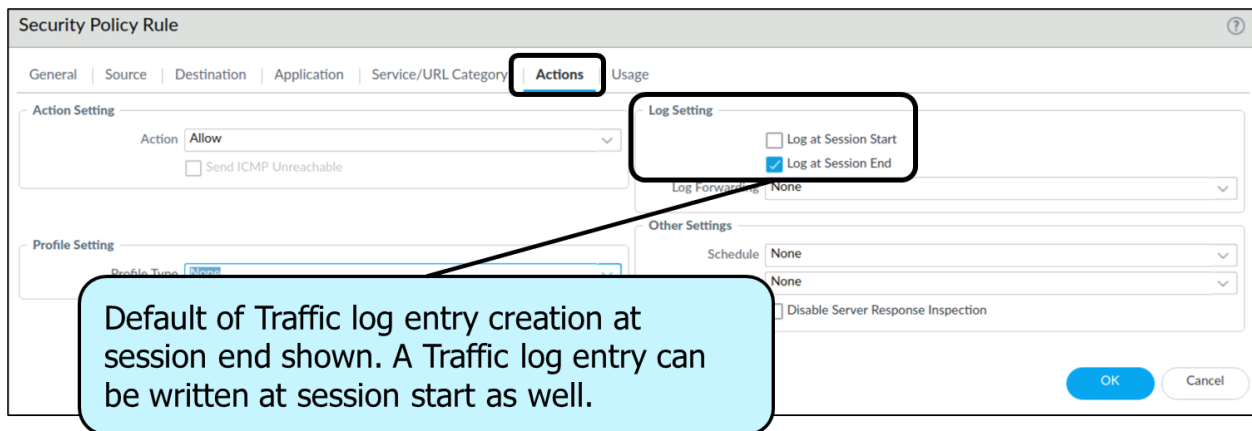
General	Source	Destination
Session ID 48755	Source User	Destination User
Action allow	Source 192.168.1.20	Destination 4.2.2.2
Action Source from-policy	Source DAG	Destination DAG
Host ID	Country 192.168.0.0-192.168.255.255	Country United States
Application dns	Port 52586	Port 53
Rule Users_to_Internet	Zone Users	Zone Internet
Rule UUID 7a2d1bdf-f078-430c-82f7-ec539e65aa41	Interface ethernet1/2	Interface ethernet1/1
Session End Reason aged-out	NAT IP	NAT IP 4.2.2.2
Category any	NAT Port 40457	NAT Port 53

#### Log Entry Detail

Log entry details include many types of information for troubleshooting: the Security action, the firewall policy allowing the traffic through, the assigned App-ID, the zones, and the ingress and

egress interfaces. Log entries also include NAT details and flags for other handling details. Examine this data to get valuable insight into the firewall's processing of traffic from both connectivity and security processing views.

This data is typically written at session end, but logging settings can specify that log entries be created at session initiation time. This practice drives more log volume, but it can provide critical data in certain situations. Configure the Log at Session Start option temporarily during troubleshooting to provide more information and gain insight, as shown in the following image:



You can also display open sessions by using the **Monitor > Session** Browser display, as shown here:

START TIME	FROM ZONE	TO ZONE	SOURCE	DESTINATION	FROM PORT	TO PORT	PROTOCOL	APPLICATION	RULE	INGRESS I/F	EGRESS I/F
06/27 00:20:10	Users_Net	Internet	192.168.1.254	130.211.8.196	47542	443	6	paloalto-dns-security	Allow-Internet-Access	ethernet1/2	ethernet1/1
06/27 00:43:24	Extranet	Extranet	192.168.50.1	192.168.50.53	34273	53	17	dns	vsys1+intr...default	ethernet1/3	ethernet1/3
			<b>Detail</b>			<b>Flow 1</b>			<b>Flow 2</b>		
			Session ID	2147				Direction	s2c		
			Timeout	30	From Zone	Extranet			Extranet		
			Time To Live	8	Source	192.168.50.1			192.168.50.53		
			Virtual System	vsys1	Destination	192.168.50.53			192.168.50.1		
			Application	dns	From Port	34273			53		
			Protocol	17	To Port	53			34273		
			Security Rule	vsys1+intrazone-default	From User	unknown			unknown		
			URL Category	any	To User	unknown			unknown		
			QoS Rule	N/A	State	ACTIVE			ACTIVE		
			QoS Class	4	Type	FLOW			FLOW		
			Created by Syn Cookie	False							
			To Host Session	True							
			Traverse Tunnel	False							
			Captive Portal	False							
			Session End Log	True							
			Session In Ager	True							
			Session From HA	False							
			End Reason	unknown							
06/27 00:43:24	Extranet	Extranet	192.168.50.1	192.168.50.53	48577	53	17	dns	vsys1+intr...default	ethernet1/3	ethernet1/3
06/27 00:43:24	Extranet	Extranet	192.168.50.1	192.168.50.53	43684	53	17	dns	vsys1+intr...default	ethernet1/3	ethernet1/3
06/27 00:43:24	Extranet	Extranet	192.168.50.1	192.168.50.53	50226	53	17	dns	vsys1+intr...default	ethernet1/3	ethernet1/3
06/27 00:43:24	Extranet	Extranet	192.168.50.1	192.168.50.53	45960	53	17	dns	vsys1+intr...default	ethernet1/3	ethernet1/3
06/27 00:22:37	Users_Net	Internet	192.168.1.254	65.154.226.123	58915	443	6	pan-db-cloud	Allow-Internet-Access	ethernet1/2	ethernet1/1

You can use the **Clear** check box at the end of a session summary line to end the session immediately, which often generates the desired log entry.

The CLI **show** commands can also help with troubleshooting. The web interface traffic capture and the CLI **pcap** and **debug** functions give greater visibility into the system-level operations.

Connectivity issues often arise from unexpected traffic-forwarding decisions. You can view forwarding decisions after you display the Layer 3 routing and forwarding tables in the web interface, as shown in the following figure:



You can see the specific virtual router's routing and forwarding tables by clicking the More Runtime Stats link.

Note that PBF policy rules can override routing decisions and must be considered when you troubleshoot connectivity. The routing and forwarding tables mentioned do not show the effects of existing PBF policy rules. PBF troubleshooting is best done on the CLI; **show** commands can display existing PBF policies and whether they are active. The **test pbf-policy-match** command shows the application of existing PBF policies on modeled traffic.

### 6.2.3 Aggregate interfaces, LACP

An aggregate interface group uses IEEE 802.1AX link aggregation to combine multiple Ethernet interfaces into a single virtual interface that connects the firewall to another network device or firewall. An aggregate group increases the bandwidth between peers by load balancing traffic across the combined interfaces. It also provides redundancy; when one interface fails, the remaining interfaces continue supporting traffic.

By default, interface failure detection is automatic only at the physical layer between directly connected peers. However, if you enable the Link Aggregation Control Protocol (LACP), failure detection is automatic at the physical and data link layers, regardless of whether the peers are directly connected. The LACP also enables automatic failover to the standby interfaces if you configured hot spares. All the Palo Alto Networks firewalls, except the VM-Series models, support aggregate groups. The Product Selection tool indicates the number of aggregate groups that each firewall supports. Each aggregate group can have up to eight interfaces.

### 6.2.4 Counters

Counters are a very useful set of indicators for the processes, packet flows, and sessions on the PA firewall and can be used to troubleshoot various scenarios.

Global counters collect the PAN-OS global counter values that are useful for troubleshooting system issues, such as packets sent and received, sessions allocated and freed, as well as packets dropped, received, and transmitted.

## 6.2.5 Tagging

Tags allow you to group objects by using keywords or phrases. You can apply tags to address objects, address groups (static and dynamic), applications, zones, services, service groups, and policy rules. You can also use an SD-WAN Interface profile to apply a link tag to an Ethernet interface. You can use tags to sort or filter objects and to visually distinguish objects by color. When you apply a color to a tag, the **Policy** tab displays the object with a background color.

You must create a tag before you group rules using that tag. After you assign grouped rules by a tag, **View Rulebase as Groups** to see a visual representation of the policy rulebase based on the assigned tags. While viewing your rulebase as groups, the policy order and priority is maintained. In this view, select the group tag to view all of the rules grouped by that tag.

## 6.2.6 References

- How to Troubleshoot the Physical Port Flap or Link Down Issue, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000PNcB>
- Monitor Transceivers, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/monitor-transceivers>
- How to Troubleshoot Using Counters via the CLI, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIXOCA0>
- Tags, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/objects/objects-tags>
- Global Counters, <https://docs.paloaltonetworks.com/pan-os/u-v/pan-os-device-telemetry-metrics-reference/device-health-and-performance/metric-dt-dhp-17>

## 6.3 Troubleshoot Decryption

### 6.3.1 Inbound decryption

Troubleshooting tools provide enhanced visibility into TLS traffic so you can monitor the decryption deployment. The tools enable you to diagnose and resolve decryption issues quickly and easily, tighten weaknesses in decryption deployment, and fix decryption issues to improve the security posture. For example, you can:

- Identify traffic that causes decryption failures by Service Name Identification (SNI) and application.
- Identify traffic that uses weak protocols and algorithms.
- Examine successful and unsuccessful decryption activity in the network.
- View detailed information about individual sessions.
- Profile decryption usage and patterns.
- Monitor detailed decryption statistics and information about adoption, failures, versions, algorithms, etc.

The following tools provide full visibility into the TLS handshake and help you troubleshoot and monitor the decryption deployment:

- **ACC > SSL Activity** — The five ACC widgets on this tab (introduced in PAN-OS 10.0) provide details about any successful and unsuccessful decryption activity in your network, including decryption failures, TLS versions, key exchanges, and the amount and type of decrypted and undecrypted traffic.
- **Monitor > Logs > Decryption** — The Decryption Log (introduced in PAN-OS 10.0) provides comprehensive information about individual sessions that match a Decryption policy, use a No Decryption policy for traffic you don't decrypt, and GlobalProtect sessions when you enable Decryption logging in GlobalProtect Portal or GlobalProtect Gateways configuration. Select which columns to display to view information, such as application, SNI, Decryption Policy Name, error index, TLS version, key exchange version, encryption algorithm, certificate key types, and many other characteristics. Filter the information in columns to identify traffic that uses particular TLS versions and algorithms, particular errors, or any other characteristics you want to investigate. By default, Decryption policies only log unsuccessful TLS handshakes. If you have the available log storage, configure Decryption policies to also log successful TLS handshakes in order to gain visibility into those decrypted sessions.
- **Local Decryption Exclusion Cache** — There are two constructs for sites that break decryption for technical reasons, such as client authentication or pinned certificates, and therefore need to be excluded from decryption—the SSL Decryption Exclusion List and the Local Decryption Exclusion Cache. The SSL Decryption Exclusion List contains the servers that Palo Alto Networks has identified to break decryption technically. Content updates keep the list up-to-date, and you can add servers to the list manually. The Local Decryption Exclusion Cache automatically adds the servers that local users encounter that break decryption for technical reasons and excludes them from decryption, provided the Decryption profile applied to the traffic allows unsupported modes (if unsupported modes are blocked, then the traffic is blocked instead of being added to the local cache).
- **Custom Report Templates for Decryption** — You can create custom reports by using four predefined templates that summarize decryption activity.

### 6.3.2 SSL forward proxy

When you configure the firewall to decrypt SSL traffic going to external sites, it functions as an SSL forward proxy. Use an SSL forward proxy decryption policy to decrypt and inspect SSL/TLS traffic from internal users to the web. SSL forward proxy decryption prevents malware, concealed as SSL encrypted traffic, from being introduced into the network by decrypting the traffic so that the firewall can apply decryption profiles and Security policies and profiles to the traffic.

In SSL forward proxy decryption, the firewall is a man in the middle (MITM) between the internal client and the external server. The firewall uses certificates to transparently represent the client to the server and the server to the client, so that the client believes it is communicating directly with the server (even though the client session is with the firewall) and the server believes it is communicating directly with the client (even though the server session is also with the firewall). The firewall uses certificates to establish itself as a trusted third party (man in the middle) for the client-server session.

### 6.3.3 SSH proxy

In an SSH proxy configuration, the firewall resides between a client and a server. SSH proxy enables the firewall to decrypt inbound and outbound SSH connections and ensures that attackers don't use SSH to tunnel unwanted applications and content. SSH decryption does not require certificates and the firewall automatically generates the key used for SSH decryption when the firewall boots up. During the bootup process, the firewall checks for an existing key. If not found, the firewall generates a key. The firewall uses the key to decrypt the SSH sessions for all of the virtual systems configured on the firewall and all the SSH v2 sessions.

SSH allows tunneling, which can hide malicious traffic from decryption. The firewall can't decrypt traffic inside an SSH tunnel. You can block all the SSH tunnel traffic by configuring a Security policy rule for the application **ssh-tunnel** with the **Action** set to **Deny** (along with a Security policy rule to allow traffic from the **ssh** application).

SSH tunneling sessions can tunnel X11 Windows and TCP packets. One SSH connection might contain multiple channels. When you apply an SSH Decryption profile to traffic, for each channel in the connection, the firewall examines the App-ID of the traffic and identifies the channel type. The channel type can be one of the following:

- session
- X11
- forwarded-tcpip
- direct-tcpip

When the channel type is *session*, the firewall identifies the traffic as allowed SSH traffic, such as SFTP or SCP. When the channel type is *X11*, *forwarded-tcpip*, or *direct-tcpip*, the firewall identifies the traffic as SSH tunneling traffic and blocks it.

The SSH Proxy Decryption profile (**Objects > Decryption Profile > SSH Proxy**) controls the session mode checks and failure checks for the SSH traffic defined in the SSH Proxy Decryption policies to which you attach the profile. The following figure shows the general best practice recommendations for the SSH Proxy Decryption profile settings, but the settings also depend on the organization's security compliance rules and local laws and regulations.

The screenshot shows the 'Decryption Profile' configuration window. The 'Name' field is 'best-practice-ssl-decryption'. The 'SSH Proxy' tab is selected. Under 'Unsupported Mode Checks', two options are checked: 'Block sessions with unsupported versions' and 'Block sessions with unsupported algorithms'. Under 'Failure Checks', two options are unchecked: 'Block sessions on SSH errors' and 'Block sessions if resources not available'. A note at the bottom states: 'Note: For unsupported modes and failures, the session information is cached for 12 hours, so future sessions between the same host and server pair are not decrypted. Check boxes to block those sessions instead.' The 'OK' button is highlighted in blue.



## Unsupported Mode Checks.

The firewall supports SSHv2. If you don't block sessions with unsupported modes, users receive a warning message if they connect with potentially unsafe servers; they can click through that message and reach the potentially dangerous site. Blocking these sessions, as described below, protects you from the servers that use weak, risky protocol versions and algorithms:

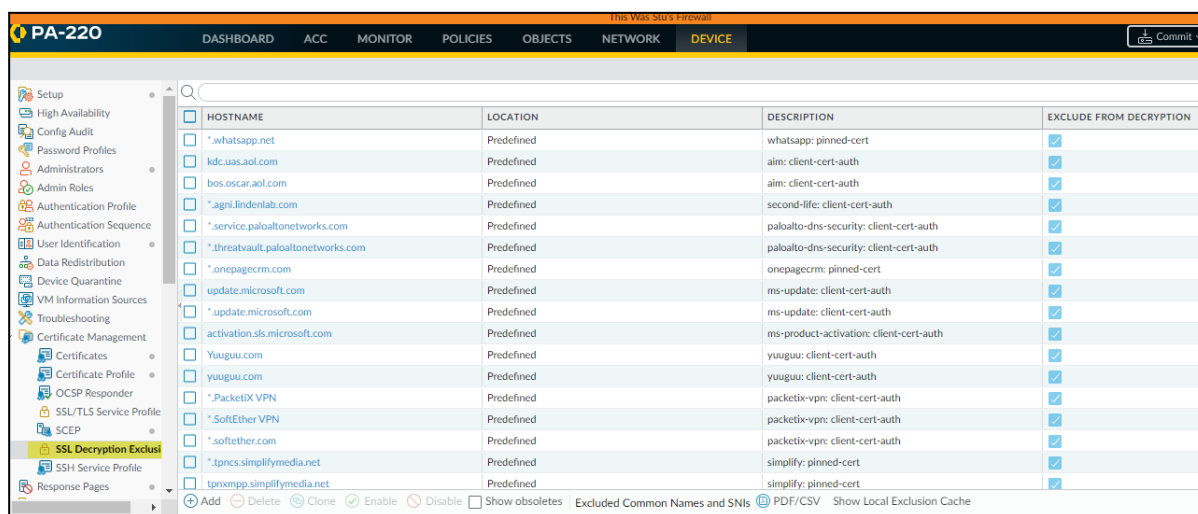
- **Block sessions with unsupported versions** — The firewall has a set of predefined supported versions. Checking this box blocks traffic with weak versions. Always check this box to block sessions with the weak protocol versions to reduce the attack surface.
- **Block sessions with unsupported algorithms** — The firewall has a set of predefined supported algorithms. Checking this box blocks traffic with weak algorithms. Always check this box to block sessions with unsupported algorithms to reduce the attack surface.

### 6.3.4 Identify what cannot be decrypted and configure exclusions and bypasses

The firewall provides a predefined SSL Decryption Exclusion list to exclude commonly used sites from decryption, which break decryption because of technical reasons such as pinned certificates and mutual authentication. The predefined decryption exclusions are enabled, by default, and Palo Alto Networks delivers new and updated predefined decryption exclusions to the firewall as part of the Applications and Threats content update (or the Applications content update if you do not have a Threat Prevention license). The firewall does not decrypt traffic that matches the predefined exclusions and allows the encrypted traffic based on the Security policy that governs that traffic. However, the firewall can't inspect the encrypted traffic or enforce a Security policy on it.

The traffic of the sites on the SSL Decryption Exclusion list remains encrypted; therefore, the firewall does not inspect or provide further security enforcement of the traffic. You can disable a predefined exclusion. For example, you can choose to disable predefined exclusions to enforce a strict security policy that allows only those applications and services that the firewall can inspect and on which the firewall can enforce the Security policy. However, the firewall blocks those sites whose applications and services break decryption technically if they are not enabled on the SSL Decryption Exclusion list.

You can view and manage all of the Palo Alto Networks predefined SSL decryption exclusions directly on the firewall (**Device > Certificate Management > SSL Decryption Exclusions**).



The screenshot shows the Palo Alto Networks firewall management interface. The top navigation bar includes 'PA-220', 'DASHBOARD', 'ACC', 'MONITOR', 'POLICIES', 'OBJECTS', 'NETWORK', and 'DEVICE'. The left sidebar shows a navigation menu with 'SSL Decryption Exclusions' selected. The main content area displays a table of predefined exclusions.

HOSTNAME	LOCATION	DESCRIPTION	EXCLUDE FROM DECRYPTION
<input type="checkbox"/> *.whatsapp.net	Predefined	whatsapp: pinned-cert	<input checked="" type="checkbox"/>
<input type="checkbox"/> kdc.usas.aol.com	Predefined	aim: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> bos.oscar.aol.com	Predefined	aim: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.agnii.lindenlab.com	Predefined	second-life: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.service.paloaltonetworks.com	Predefined	paloalto-dns-security: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.threatvault.paloaltonetworks.com	Predefined	paloalto-dns-security: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.onepagecrm.com	Predefined	onepagecrm: pinned-cert	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.update.microsoft.com	Predefined	ms-update: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.update.microsoft.com	Predefined	ms-update: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.activation.sls.microsoft.com	Predefined	ms-product-activation: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> Yuuguu.com	Predefined	yuuguu: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> yuuguu.com	Predefined	yuuguu: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.PacketIX VPN	Predefined	packetix-vpn: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.SoftEther VPN	Predefined	packetix-vpn: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.softether.com	Predefined	packetix-vpn: client-cert-auth	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.tpncs.simpliflymedia.net	Predefined	simplify: pinned-cert	<input checked="" type="checkbox"/>
<input type="checkbox"/> *.tpnxmp.simpliflymedia.net	Predefined	simplify: pinned-cert	<input checked="" type="checkbox"/>



### 6.3.5 Certificates

Below are general verification checks for certificates. The certificates are categorized as either self-generated or purchased.

#### Self-Generated Certificates

- Check Certificate Details
- Check Validity Date
- Check CA Distribution To Client Browsers
- Check Intermediate CA (signed by root CA or another Intermediate CA)
- Check Certificate Chain

#### Purchased Certificates

- Check Issuing Authority
- Check Validity Date
- Check CA Distribution to Client Browsers
- Check Intermediate CA (signed by root CA)
- Check Certificate Chain

### 6.3.6 References

- Troubleshoot and Monitor Decryption,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/troubleshoot-and-monitor-decryption>
- SSL Forward Proxy,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/decryption-concepts/ssl-forward-proxy>
- SSH Proxy,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/decryption-concepts/ssh-proxy>
- Troubleshooting SSL Certificates  
<https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000ClgbCAC>
- Palo Alto Networks Predefined Decryption Exclusions,  
<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/decryption/decryption-exclusions/palo-alto-networks-predefined-decryption-exclusions>

## 6.4 Troubleshoot routing

The NGFW uses several methods and configurations to route traffic, as are described in the following sections.

### 6.4.1 Dynamic routing

- **Routing Information Protocol (RIP)** should be troubleshoot with the following considerations in mind:
  - Interface shut: A network on an interface that is in shutdown will not be advertised.
  - Passive interface: An interface that has been configured as passive will not send any RIP updates.
  - Version mismatch: RIP has two versions, both routers should use the same version.
  - Max hop count: When the hop count is 16, the network is considered unreachable.
  - Route Filtering: Filters might prevent RIP updates from being sent or received.
  - Authentication: Both RIP routers should have the same authentication parameters.
  - Split horizon: Networks that are learned on an interface are not advertised out of the same interface.
  - Auto-summarization: Causes issues with non-contiguous networks.
- **Open Shortest Path First (OSPF)** has the following elements to be considered when troubleshooting if your adjacency with the Peer is stuck in INIT State:
  - Check the firewall to verify if outbound OSPF packets are being dropped.
  - If the firewall is seen transmitting OSPF Hellos out, then you should check the peer router to see if it is receiving these hellos on its interface.
  - If the packets are not seen on the peer, then the switch could be dropping (maybe vwire device in between), and this needs to be checked.
  - If the packets are being received on the peer, then you should check Peer's logs to see if it is dropping them due to any reason (normally, parameters should be matching since you accepted Peer's Hello).

If you see that the state is stuck in EXCHANGE or LOADING State, it means that either LSA Headers are not getting exchanged or LSA Updates are not getting exchanged.

- You should make sure MTU matches on the OSPF Peers and on the device (switch/transparent firewall) in between.
- You should make sure bidirectional communication is allowed all the way, i.e. on both PA and OSPF Peer and also on the device in between. There should not be any ACLs blocking any such packets.
- Fragmentation could be an issue since LSA updates are large and the device in between (or the firewall interface itself) has a smaller MTU, so we should make sure there are no settings to block fragments.
- On Palo Alto debugs, both pcaps and global counters can again help to verify if we are sending out and receiving packets, but simultaneous captures on the peer are also important.

- **Border Gateway Protocol (BGP)** - Routing level issues include routing protocol configurations issues, BGP neighbor establishment issues, static routing or misconfigured static routes, or issues with prefix learning or advertisement.
  - For BGP session establishment related issues, check if there is an incorrect AS# or global parameters, an incorrect BGP peer IP, or check if BGP multi-hop is required.
  - For BGP peer type issues (data center only), check if the right peer type is selected, core, edge, or classic peer type. The edge peer only learns prefixes.
  - For static routing related issues, check the configuration for administrative distance, and next hop.
  - For prefix learning and advertisement issues, check the route map configurations on firewall and BGP peer devices, check for interactions with other routing protocols in the enterprise network, and check for split or no-split prefix scenarios.

### 6.4.2 Redistribution profiles

Route redistribution on the firewall is the process of making the routes that the firewall learned from one routing protocol (or a static or connected route) available to a different routing protocol. This increases the accessibility of network traffic. Without route redistribution, a router or virtual router advertises and shares routes only with other routers that run the same routing protocol. You can redistribute IPv4 or IPv6 BGP, connected, or static routes into the OSPF RIB and redistribute OSPFv3, connected, or static routes into the BGP RIB.

You can make specific networks, once available only by manual static route configuration on specific routers, available to BGP autonomous systems or OSPF areas. You can also advertise locally connected routes, such as routes to a private lab network, into BGP autonomous systems or OSPF areas.

- Give users on the internal OSPFv3 network access to BGP so they can access the devices on the internet. In this case, you would redistribute BGP routes into the OSPFv3 RIB.
- Give external users access to some parts of the internal network, so you make internal OSPFv3 networks available through BGP by redistributing OSPFv3 routes into the BGP RIB.

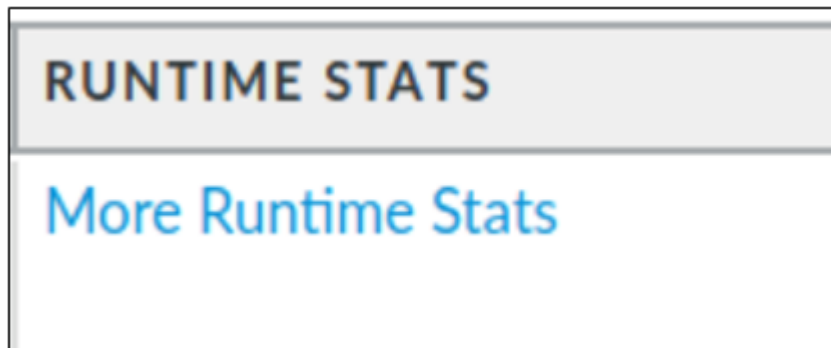
### 6.4.3 Static routes

Common issues for static routes include: invalid route monitoring profile information; incorrect subnet masks; incorrect next hops.

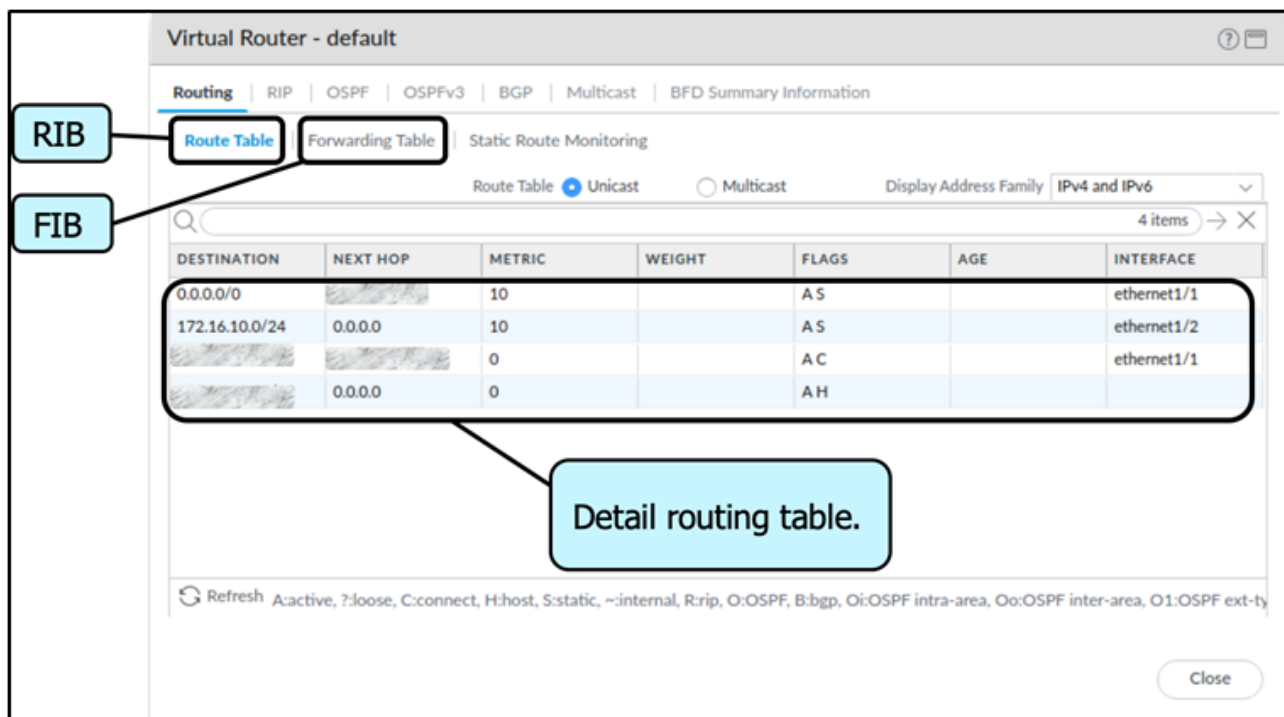
## 6.4.4 Route monitoring

### Routing Troubleshooting

The routing decisions made by a virtual router can be diagnosed easily. A virtual router maintains an RIB and an FIB, which can be displayed in the management web interface by using the **Runtime Stats** link displayed on the virtual router summary line:



Click the **More Runtime Stats** link to access the routing table (RIB) and the forwarding table (FIB) with additional displays that contain the status of any enabled dynamic routing protocols.



The screenshot shows the 'Virtual Router - default' configuration page. The 'Routing' tab is active, and the 'Route Table' sub-tab is selected. The interface displays a table of routing entries. A callout box labeled 'RIB' points to the 'Route Table' tab, and another callout box labeled 'FIB' points to the 'Forwarding Table' tab. A third callout box labeled 'Detail routing table.' points to the table of entries.

DESTINATION	NEXT HOP	METRIC	WEIGHT	FLAGS	AGE	INTERFACE
0.0.0.0/0		10		A S		ethernet1/1
172.16.10.0/24	0.0.0.0	10		A S		ethernet1/2
		0		A C		ethernet1/1
	0.0.0.0	0		A H		

### Troubleshooting Routing

The CLI has advanced troubleshooting of routing functions. Output from the debug **routing ...** command provides insight into router processing, including advanced debugging logs and routing-specific packet captures.

## 6.4.5 Policy-based forwarding

Often, the underlying cause of PBF issues has to do with conflicts with static or dynamic routes. Keep in mind that BPF rules take precedence. Troubleshooting PBF rules requires reviewing the order in which PBF rules are processed, evaluating what traffic would slip through to other defined routes, and looking at session browser details for connected sessions to see what next hop they are taking and which rules are being applied.

## 6.4.6 Multicast routing

The following steps should be reviewed when attempting to verify functionality of multicast:

### Static Rendezvous Point (RP), IGMP and Security Policy Configuration Steps

1. From the WebGUI, go to **Network > Virtual Routers > Multicast**
  - Enable Multicast
  - Select Static RP, RP Interface and IP of one of the Zones that will participate in Multicast
  - Add designated Multicast Group IP
2. From the Interfaces Tab, add the Multicast / Interface Group and include all interfaces participating in multicast.
  - Enable IGMP and PIM on all interfaces  
**Note:** The firewall has to have PIM enabled, otherwise multicast routing would fail
3. Configure Security policy to allow multicast traffic
  - Include all multicast zones for Source Zone
  - Use predefined Multicast Zone for Destination Zone  
**Note:** Do not create this Zone it is a predefined Zone.
4. Commit the configuration

### For Testing and Verification of Multicast Traffic

1. Verify multicast IGMP membership. All interfaces with current IGMP traffic should be shown:

```
admin@61-PA-500> show routing multicast igmp membership
VIRTUAL ROUTER: default
-----
interface      group          source         up time      expiry      filter mode  excl mode  expiry  v1 host timer  v2 host timer  last reporter
-----
ethernet1/4    234.6.6.6     0.0.0.0        29.82       238.70     exclude     238.70     0.00     238.70         192.168.61.216
ethernet1/4    239.255.255.250 0.0.0.0        27.81       232.19     exclude     232.19     0.00     232.19         192.168.61.216
ethernet1/7    234.6.6.6     0.0.0.0        25.63       235.73     exclude     235.73     0.00     235.73         10.10.10.2
ethernet1/7    239.255.255.250 0.0.0.0        163.74      241.73     exclude     241.73     0.00     241.73         10.10.10.2
```

2. Run the CLI command: **> show routing multicast igmp statistics**

Interfaces that are currently processing multicast traffic should have a positive number of joins and queries sent.

**Note:** In the example below, the Interface Name - ethernet1/3 there aren't any "number of joins" because there were no clients that were requesting to join on that network.

```
VIRTUAL ROUTER: default

interface name: ethernet1/3

total groups: 0
total source-group pairs: 0
wrong version queries: 0
number of joins: 0
failed joins: 0
general queries sent: 614
specific queries sent: 0
total received messages: 0
received v1 messages: 0
received v2 messages: 0
received v3 messages: 0
received invalid messages: 0
peak number of groups: 0

interface name: ethernet1/4

total groups: 1
total source-group pairs: 0
wrong version queries: 0
number of joins: 33
failed joins: 0
general queries sent: 615
specific queries sent: 6
total received messages: 5926
received v1 messages: 0
received v2 messages: 5921
received v3 messages: 5
received invalid messages: 5773
peak number of groups: 2

interface name: ethernet1/7

total groups: 0
total source-group pairs: 0
wrong version queries: 0
number of joins: 2
failed joins: 0
general queries sent: 6
specific queries sent: 0
total received messages: 11
received v1 messages: 0
received v2 messages: 9
received v3 messages: 2
received invalid messages: 0
peak number of groups: 2
```

- Run the following CLI command: **> show routing multicast pim state**  
Verify the Sender IP. In this example, it is 192.168.61.216

```
admin@61-PA-500> show routing multicast pim state
VIRTUAL ROUTER: default
(*, G):
-----
group      RP      up time  upstream join st  upstream join timer  RPF interface  RPF next hop
-----
234.6.6.6  192.168.61.1  93.99   Joined            0.00                0              0.0.0.0
(*, G, I):
-----
group      interface  local membership  join/prune st  prune pending timer  join expiry timer  assert st  assert timer  assert winner addr  assert winner metric
-----
234.6.6.6  ethernet1/4  yes              NoInfo      0.00                0.00                Winner    87.19         0.0.0.0             0
234.6.6.6  ethernet1/7  yes              NoInfo      0.00                0.00                NoInfo    0.00         0.0.0.0             0
(S, G):
-----
group      source      up time  upstream nbr  upstream join st  upstream join timer  RPF next hop  DR reg st  DR req stop timer  SPT
-----
234.6.6.6  192.168.61.216  92.11   0.0.0.0      Joined            0.00                192.168.61.216  Join       0.00                yes
(S, G, rpt):
-----
group      source      up time  upstream prune st  upstream override timer
-----
234.6.6.6  192.168.61.216  92.11   NotPrune      0
```

- Run the following CLI Command: **> show routing multicast fib**  
Verify the multicast group that includes designated multicast interfaces

```
admin@61-PA-500> show routing multicast fib
VIRTUAL ROUTER: default
maximum of entries: 3
-----
group      source      flags  incoming  outgoing
-----
234.6.6.6  0.0.0.0     0      0          ethernet1/4
                ethernet1/7
239.255.255.250  10.66.24.222  96      0          0
234.6.6.6  192.168.61.216  0      ethernet1/4  ethernet1/7
```

### 6.4.7 Service routes

Suggestions for verifying service route functionality include:

**When there are no services routes configured, the table is empty.**

```
admin@PAN-5050-243(active)> debug dataplane internal vif route 250
```

```
admin@PAN-5050-243(active)>
```

**After committing the configuration, routing table has been populated:**

```
admin@PAN-5050-243(active)> debug dataplane internal vif route 250
```

```
193.190.138.68 via 172.16.31.244 dev eth3.1 src 172.16.31.244
<>199.167.52.13 via 172.16.31.244 dev eth3.1 src 172.16.31.244
195.200.224.66 via 172.16.31.244 dev eth3.1 src 172.16.31.244
85.234.197.3 via 172.16.31.244 dev eth3.1 src 172.16.31.244
192.168.200.99 via 172.16.31.244 dev eth3.1 src 172.16.31.244
10.192.16.98 via 172.16.31.244 dev eth3.1 src 172.16.31.244
85.234.197.4 via 172.16.31.244 dev eth3.1 src 172.16.31.244
8.8.8.8 via 172.16.31.244 dev eth3.1 src 172.16.31.244
```

This means that these routes are active in the Management Plane.

#### 6.4.8 References

- Route Redistribution, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/route-redistribution>
- Multicast, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/network/network-routing-routing-profiles/network-routing-routing-profiles-multicast>
- Service Routes, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-networking-admin/service-routes/service-routes-overview#id69ef535a-d5b0-4c79-bb7f-1302a438e7c5>

### 6.5 General Troubleshooting

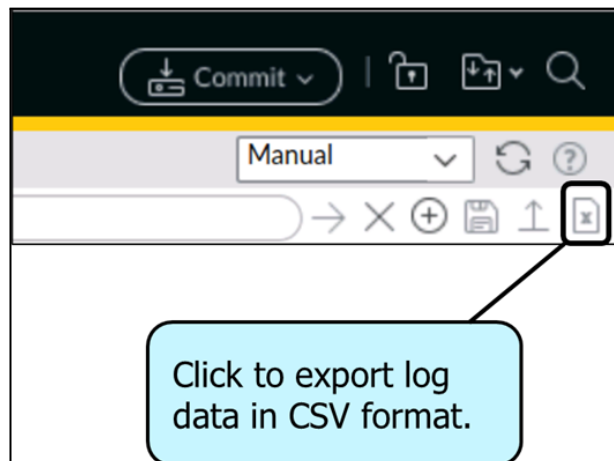
#### 6.5.1 Logs

Logging and reporting are critical components of any security network. The ability to log all the network activity in a logical, organized, and easily segmented way makes logging even more valuable. Rapid, thorough, and accurate interpretation of events is critical to security. Security practitioners often suggest that security is only as good as the visibility that it provides. These reasons contribute to the information collection and display design of Palo Alto Networks.

Log information is generally in the **Monitor** tab of the web interface. The reporting sections align with the general use of these reports. The **Log** section presents both detailed, real-time, and historical data (subject to available storage). The information is subdivided into sections that segment log data into related information. The PAN-OS software includes a Unified log that collects copies of events from the Traffic, Threat, URL Filtering, Data Filtering, WildFire Submissions, Tunnel Inspection, Authentication, and IP-Tag logs into a single location for easy parsing of related data.

Each log provides similar features, which results in an organized presentation of desired data. Displayed log data can be exported in the CSV format at any time.

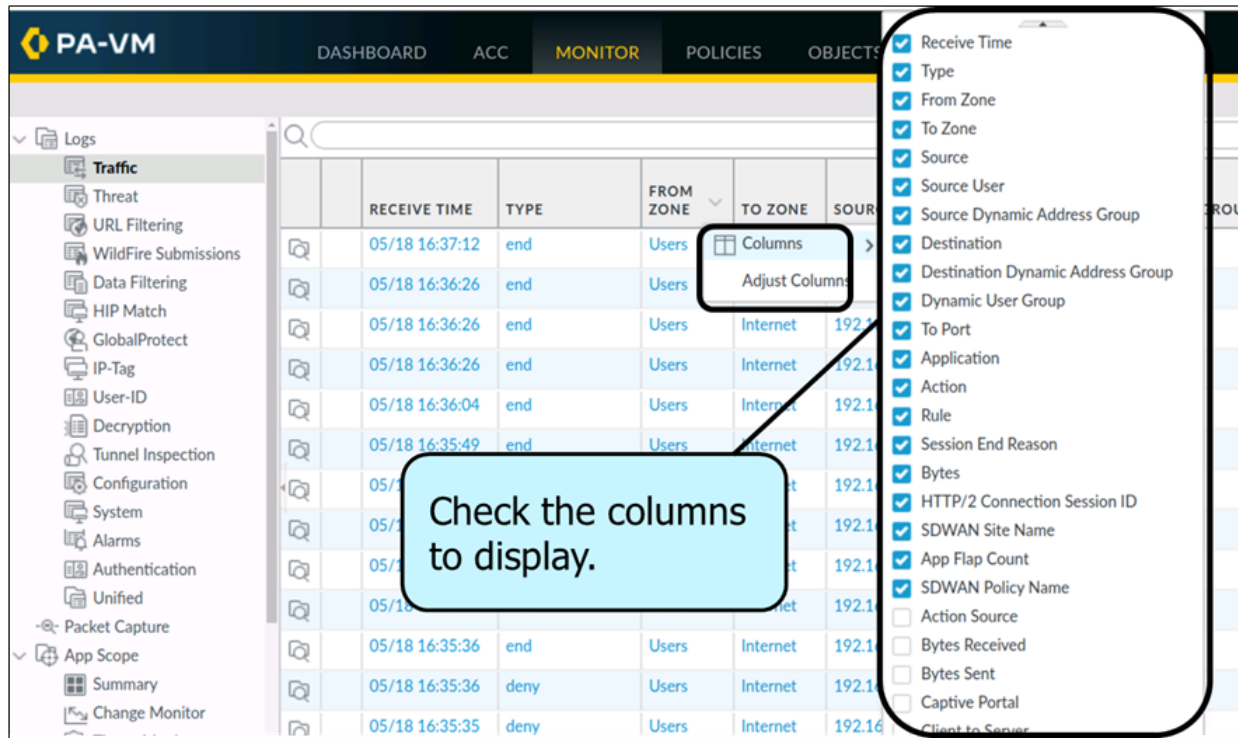
The following figure shows the CSV export option available on any detailed log display:



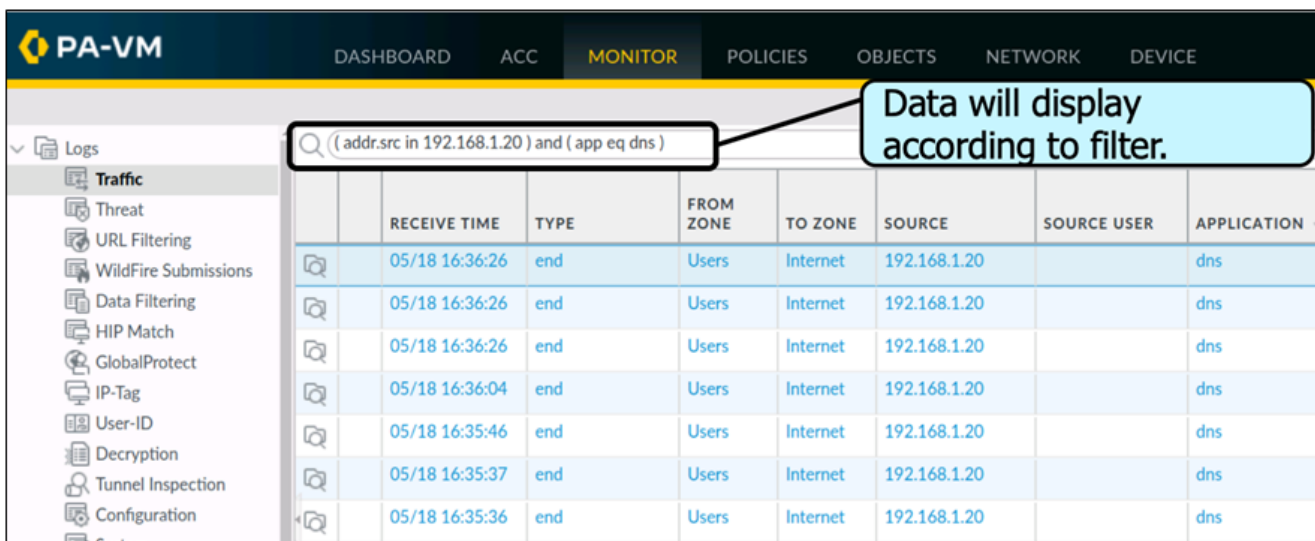


This export includes all the details for the displayed record even if it is not visible in the chosen column displays.

You can see the entries in various logs by using **Monitor > Logs**. You can configure the columns to display and their order and width:



Each log display offers a powerful filtering capability that facilitates the display of specific desired data:



Filters can be added to eliminate the display of undesired entries.

## 6.5.2 Packet capture (pcap)

All the Palo Alto Networks firewalls allow you to take packet captures (pcaps) of traffic that traverses the management interface and network interfaces on the firewall. When taking pcaps on the data plane, you may need to disable hardware offload to ensure that the firewall captures all of the traffic.

Packet capture can be very CPU-intensive and can degrade firewall performance. Use this feature only when necessary and make sure you turn it off after collecting the required packets.

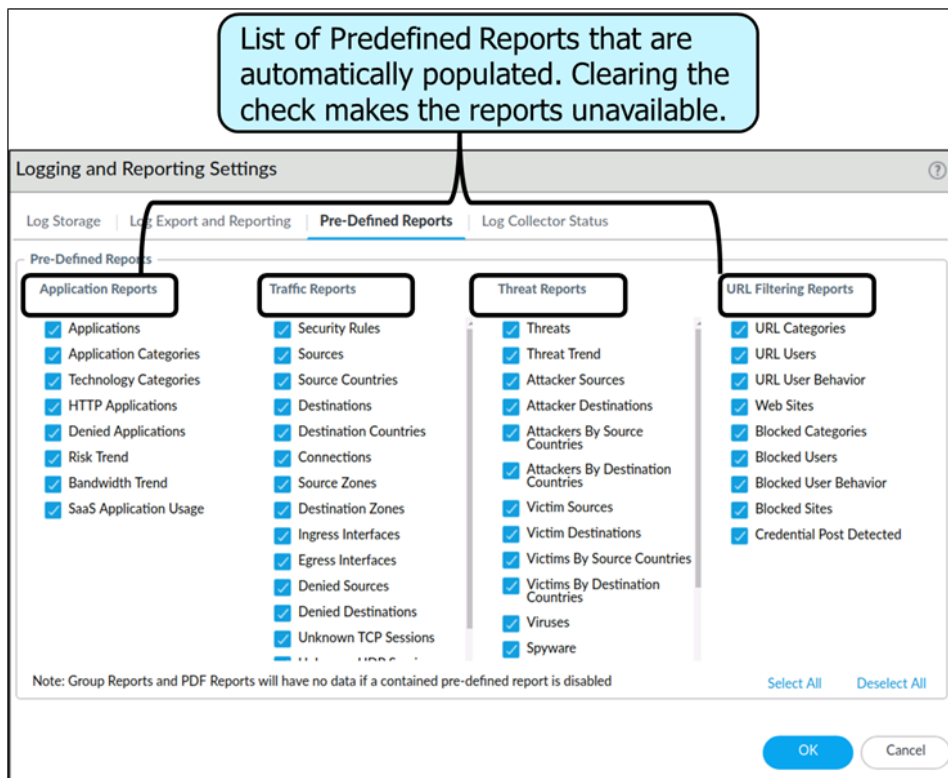
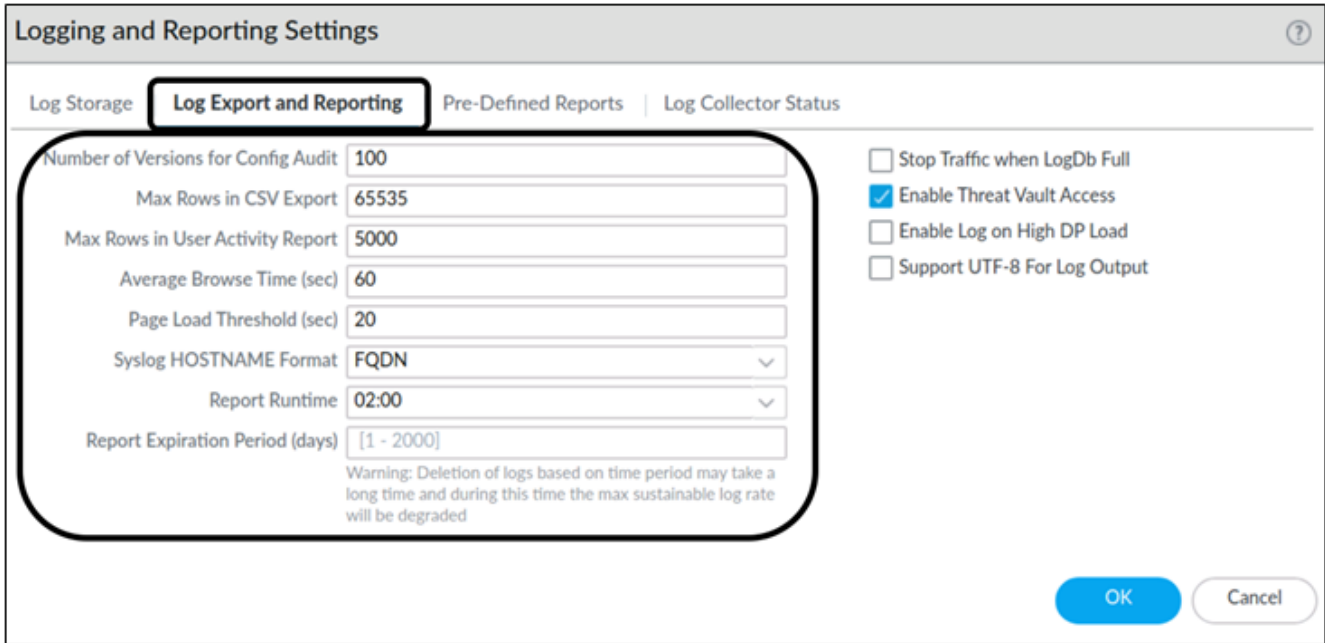
The different types of pcaps you can enable depend on what you need to do, as described:

- **Custom Packet Capture** — The firewall captures packets for all of the traffic or for specific traffic based on the filters that you define. For example, you can configure the firewall to only capture packets to and from a specific source and destination IP address or port. You can then use the packet captures for troubleshooting network-related issues or for gathering application attributes to enable writing custom application signatures or requesting for an application signature from Palo Alto Networks.
- **Threat Packet Capture** — The firewall captures packets when it detects a virus, spyware, or vulnerability. You enable this feature in the Antivirus, Anti-Spyware, and Vulnerability Protection Security profiles. A link to view or export the packet captures appears in the second column of the Threat log. These packet captures provide context around a threat to help you determine if an attack is successful or to learn more about the methods used by an attacker. You can also submit this type of pcap to Palo Alto Networks to have a threat re-analyzed if you feel it's a false-positive or false-negative.
- **Application Packet Capture** — The firewall captures packets based on a specific application and the filters that you define. A link to view or export the packet captures appears in the second column of the Traffic logs for the traffic that matches the packet capture rule.
- **Management Interface Packet Capture** — The firewall captures packets on the management interface. The packet captures are useful when troubleshooting services that traverse the interface, such as firewall management authentication to external authentication services, software and content updates, log forwarding, communication with SNMP servers, and authentication requests for GlobalProtect and Captive Portal.
- **GTP Event Packet Capture**—The firewall captures a single GTP event, such as GTP-in-GTP, end user IP spoofing, and abnormal GTP messages, to make GTP troubleshooting easier for mobile network operators.

### 6.5.3 Reports

While log data is stored in detail in log storage, a firewall summarizes new log entries and adds the results to separate the on-board reporting databases used as default sources by Application Command Center (ACC), App Scope, PDF Reports, and Custom Reports.

The scope of this summarization process can be controlled with the settings on **Device > Setup > Management > Logging and Reporting Settings**, as shown:



## PDF Reports

The PDF Reports section offers many predefined PDF reports that can be run as a group on a scheduled basis and delivered through email daily or weekly.

By default, these reports typically run once per day and summarize all the activity on the firewall. A report browser of predefined reports appears on the right. In the figure below, chosen reports display their results for the previous day's traffic. The predefined report browser shows the choices of categories and specific reports on the right:

The screenshot displays the PA-VM interface with the following components and callouts:

- 2**: Points to the "Reports" menu item in the left navigation pane.
- 3**: Points to the "Destination Countries" report in the "Traffic Reports" section of the report browser.
- 4**: Points to the date range selector in the report browser, showing "April 2020" with the 12th selected.
- Result**: Points to the table displaying traffic data for destination countries.

	DESTINATION COUNTRY	BYTES	SESSIONS
1	192.168.0.0-192.168.255.255	22.0M	152.4k
2	United States	141.0M	18.3k
3	United Kingdom	13.3k	47
4	Ireland	6.2k	1

The PDF Reports section offers other reporting tools. Custom reports can be created, stored, and run on-demand and/or on a schedule basis.

## SaaS Applications

The App-ID engine identifies SaaS applications and provides additional functionality. A dedicated SaaS Application Usage report under **Monitor > PDF Reports > SaaS Application Usage** helps organizations identify applications that store data in external locations. The App-IDs for SaaS applications contain additional data about these applications and their providers to help you make decisions, allowing them access at the organizational level.

The screenshot shows the configuration page for the application 'dropbox-base'. The page is divided into several sections:

- Name:** dropbox-base (highlighted with a black box)
- Description:** Dropbox is a file hosting service that offers cloud storage, file synchronization, personal cloud, and client software.
- Standard Ports:** tcp/17500, tcp/443, tcp/80, udp/17500
- Depends on:** google-base
- Implicitly Uses:** ssl, web-browsing
- Deny Action:** drop-reset
- Additional Information:** Website Wikipedia Google Yahoo!
- Characteristics:**
  - Evasive: yes
  - Excessive Bandwidth Use: no
  - Used by Malware: no
  - Capable of File Transfer: yes
  - Has Known Vulnerabilities: yes
  - Tunnels Other Applications: no
  - Prone to Misuse: no
  - Widely Used: yes
  - SaaS: yes
- Options:**
  - Session Timeout (seconds): 30 (Customize...)
  - TCP Timeout (seconds): 3600 (Customize...)
  - UDP Timeout (seconds): 30 (Customize...)
  - TCP Half Closed (seconds): 120 (Customize...)
  - TCP Time Wait (seconds): 15 (Customize...)
  - App-ID Enabled: yes
- Classification:**
  - Category: general-internet
  - Subcategory: file-sharing
  - Risk: 4 (Customize...)
- SaaS Characteristics:**
  - Certifications: HIPAA, PCI, SOC I, SOC II, SSAE16
  - Data Breaches: no
  - IP Based Restrictions: no
  - Poor Financial Viability: no
  - Poor Terms Of Service: no
- Tags:** Web App (with an Edit button)

A blue callout box with the text "Additional information about this SaaS application." points to the SaaS Characteristics section. A "Close" button is located at the bottom right of the page.

Palo Alto Networks firewalls include a feature within the URL Filtering engine that provides HTTP Header Insertion for certain SaaS applications that can prevent users from accessing private instances of a SaaS application while having access to the organization's sanctioned environment User/Group Activity Report.

A predefined User Activity report provides the complete application use and browsing activity reports for individuals or groups.

A PDF Summary Report includes several top five oriented reports grouped to provide a general representation of the firewall's traffic during the previous day.

App Scope reports focus on the baseline performance comparisons of firewall use. These reports provide power tools to characterize changes in detected use patterns. They were designed for ad-hoc queries more than for scheduled report output.

#### 6.5.4 References

- View and Manage Logs, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/view-and-manage-logs>
- Packet Captures, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/monitoring/take-packet-captures>

## 6.6 Troubleshoot resource protections

### 6.6.1 Zone Protection profiles

Typical troubleshooting of zone protection profiles will involve performing packet captures to verify traffic rates and content. In terms of settings, the zone to which the profile is bound will be a common configuration element to review.

### 6.6.2 DoS protections

DoS Protection is resource-intensive, so use it only for critical systems. Similar to Zone Protection profiles, DoS Protection profiles specify flood thresholds. DoS Protection policy rules determine the devices, users, zones, and services to which DoS Profiles apply. Troubleshooting these profiles will often involve verification of traffic flows and packet captures to verify packet contents and flow rates.

### 6.6.3 Packet buffer protections

The following settings should be reviewed for packet buffer protections:


- **Global Packet Buffer Protection**—The firewall monitors sessions from all of the zones (regardless of whether Packet Buffer Protection is enabled in a zone) and how those sessions utilize the packet buffer. You must configure Packet Buffer Protection globally (**Device > Setup > Session Settings**) to protect the firewall and to enable it on individual zones. When packet buffer consumption reaches the configured **Activate** percentage, the firewall uses Random Early Drop (RED) to drop packets from the offending sessions (the firewall doesn't drop complete sessions at the global level).
- **Per-Zone Packet Buffer Protection**—Enable Packet Buffer Protection on each zone (**Network > Zones**) to layer in a second level of protection. When packet buffer consumption crosses the **Activate** threshold and global protection begins to apply RED to session traffic, the **Block Hold Time** timer starts. The **Block Hold Time** is the amount of time in seconds that the offending session can continue before the firewall blocks the entire session. The offending session remains blocked until the **Block Duration** time expires.

#### 6.6.4 References

- Zone Protection, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/network/network-network-profiles/network-network-profiles-zone-protection>
- DoS Protection Profiles and Policy Rules, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/zone-protection-and-dos-protection/zone-defense/dos-protection-profiles-and-policy-rules>
- Zone Protection and DoS Protection, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/zone-protection-and-dos-protection>
- Packet Buffer Protection, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-admin/zone-protection-and-dos-protection/zone-defense/packet-buffer-protection>

#### 6.7 Troubleshoot GlobalProtect

GlobalProtect has three major components: the GlobalProtect portal, GlobalProtect gateways, and GlobalProtect client software. These components provide management functions for the GlobalProtect infrastructure.




##### 6.7.1 Portal and Gateway

Troubleshooting GlobalProtect will often entail the same steps as for troubleshooting IPSec connections, and also the steps for troubleshooting certificate issues. In addition, you should include a review of GlobalProtect logs and reports. The ACC displays a graphical view of user activity in the GlobalProtect deployment on the GlobalProtect Activity tab.



The following GlobalProtect charts are available:

- **Successful GlobalProtect Connection Activity:** Chart view of the GlobalProtect connection activity over the selected time period. Use the toggle at the top of the chart to switch between connection statistics by users, portals and gateways, and location.
- **Unsuccessful GlobalProtect Connection Activity:** Chart view of the unsuccessful GlobalProtect connection activity over the selected time period. Use the toggle at the top of the chart to switch between connection statistics by users, portals and gateways, and location. To help you identify and troubleshoot connection issues, you also can view the reasons chart or graph. For this chart, the ACC indicates the error, source user, public IP address, and other information to help you identify and quickly resolve the issue.
- **GlobalProtect Deployment Activity:** Chart view summary of your deployment. Use the toggle at the top of the chart to view the distribution of users by authentication method, GlobalProtect app version, and operating system version.



**GlobalProtect Activity** charts and graphs are interactive and support similar drill-down functionality to other ACC charts and graphs.

The GlobalProtect Host Information widget under the **Network Activity** tab displays information about the state of the hosts on which the GlobalProtect agent is running; the host system is a GlobalProtect endpoint. This information is sourced from the entries in the HIP match log that are generated when the data submitted by the GlobalProtect app matches an HIP object or HIP Profile defined on the firewall. If you do not have HIP Match logs, this widget remains blank.

### GlobalProtect Log

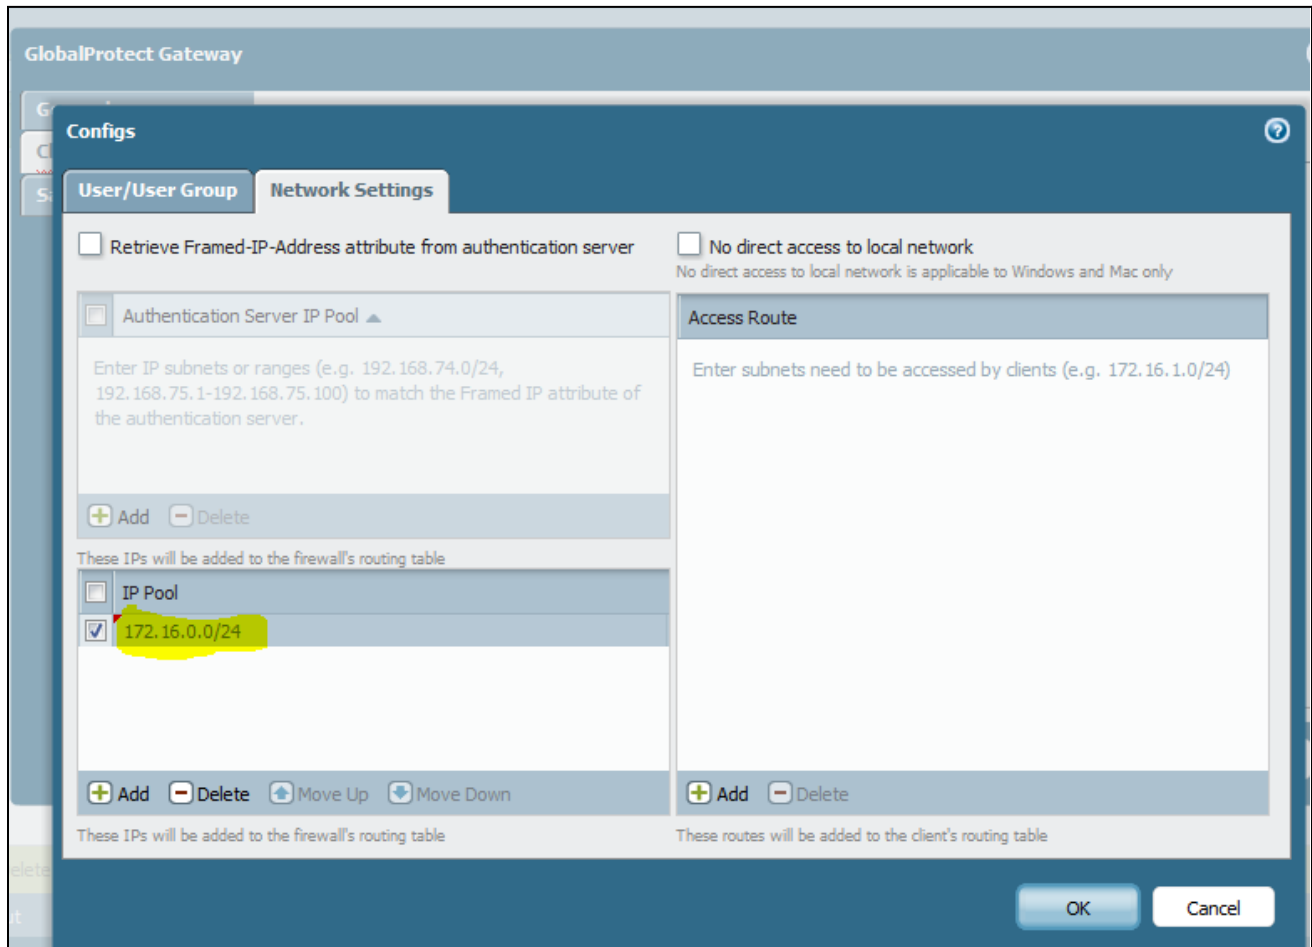
GlobalProtect logs display the following logs related to GlobalProtect:

- GlobalProtect system logs  
GlobalProtect authentication event logs in **Monitor > Logs > System**; the **Auth Method** column of the GlobalProtect logs that display the authentication method used for logins
- LSVPN/satellite events
- GlobalProtect portal and gateway logs
- Clientless VPN logs



## 6.7.2 Access to resources

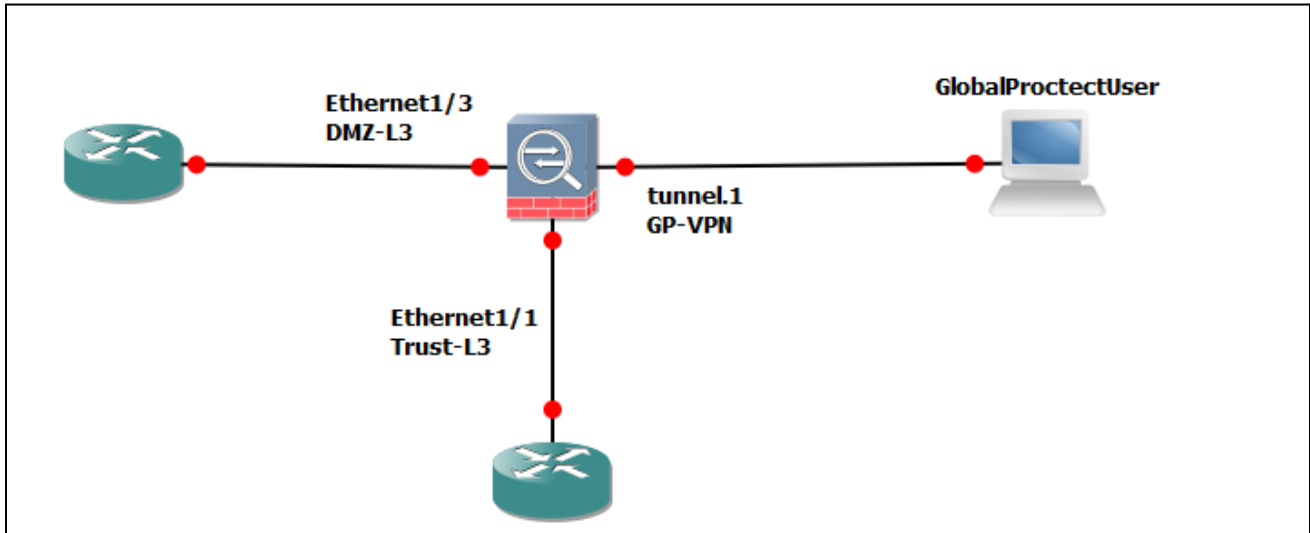
Sometimes even if the configuration is correct, GlobalProtect users are unable to access internal resources. This situation might occur because the subnet assigned to GlobalProtect is used somewhere in the network or there is a routing issue.



A workaround is to put the tunnel interface used in the GlobalProtect configuration in a different zone (GP-VPN) and do a source NAT for desired traffic. Make sure you have a Security policy to allow the traffic.

The topology:

GlobalProtect users are in the GP-VPN zone, servers are in the DMZ-L3 zone, and the internal host is in the Trust-L3 zone.



If you try to access the resources in the DMZ-L3 zone, then do a source NAT from GP-VPN to DMZ-L3.

Name	Tags	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation
1 GP-DMZ	none	GP-VPN	DMZ-L3	any	any	any	any	dynamic-ip-and-port ethernet1/3 10.50.242.57/24
2 GP-Trust-L3	none	GP-VPN	Trust-L3	any	any	any	any	dynamic-ip-and-port ethernet1/1 10.50.240.57/24

Security Policy:

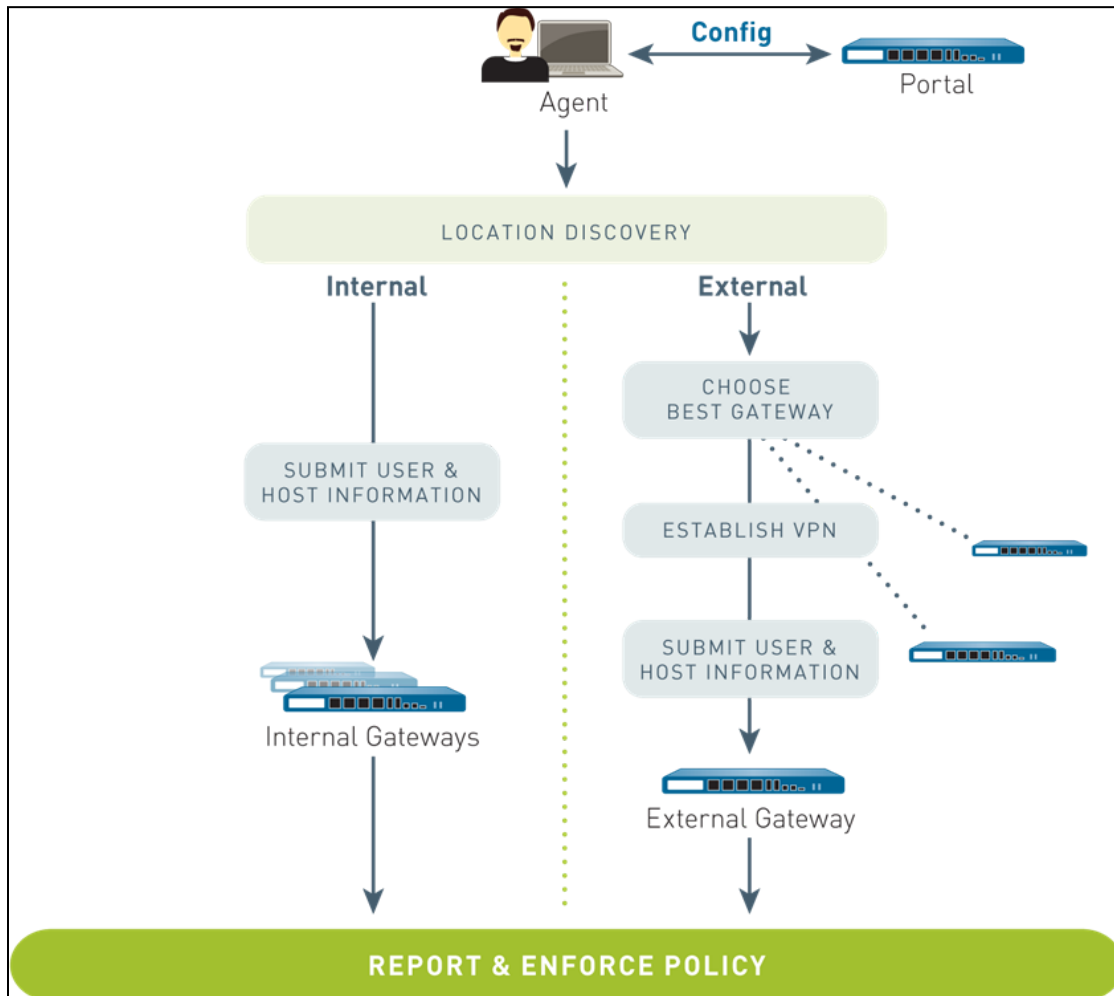
Name	Tags	Type	Zone	Address	User	HIP Profile	Zone	Address	Application	Service	URL Category	Action
1 GP-VPN To DMZ-L3 or Trust-L3	none	universal	GP-VPN	any	any	any	DMZ-L3 Trust-L3	any	any	any	any	Allow

### 6.7.3 GlobalProtect client

The GlobalProtect client software runs on end user systems and enables access to network resources via the GlobalProtect portals and gateways deployed. There are two types of GlobalProtect clients:

- **The GlobalProtect Agent** — Runs on Windows and macOS systems and is deployed from the GlobalProtect portal. You configure agent behavior—for example, which tabs the users can see—in the client configuration(s) defined on the portal.
- **The GlobalProtect App** — Runs on iOS, Android, Windows UWP, and Chromebook devices. Users must obtain the GlobalProtect app from the Apple App Store (for iOS), Google Play (for Android), Microsoft Store (for Windows UWP), or Chrome Web Store (for Chromebook).

The following diagram illustrates how the GlobalProtect portals, gateways, and agents/apps work together to enable secure access for all of your users, regardless of the devices they are using or their location.



#### 6.7.4 References:

- GlobalProtect Gateways, <https://docs.paloaltonetworks.com/globalprotect/10-1/globalprotect-admin/globalprotect-gateways/globalprotect-gateways-overview>
- GlobalProtect Users and Internal Resources, <https://knowledgebase.paloaltonetworks.com/kCSArticleDetail?id=kA10g000000ClA8>
- GlobalProtect Resource List on Configuring and Troubleshooting, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000ClfXCAS>  
Troubleshooting GlobalProtect, <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000ClkBCAS>
- GlobalProtect Client, <https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/globalprotect/device-globalprotect-client>

## 6.8 Troubleshoot policies

### 6.8.1 NAT

When troubleshooting potential NAT issues, the traffic logs and session browser are two of the most important tools. In addition, the `test nat-policy-match` command in the CLI is also of value.

### 6.8.2 Security

Typical methods for troubleshooting Security Policy rules matches include: utilizing traffic logs and the session browser to confirm which rules are applied to traffic; using the `test security-policy match` command in the CLI; and checking for rule shadows.

### 6.8.3 Decryption

Traffic logs and the session browser are helpful to verify if SSL Decryption is being applied to traffic. In addition, shadows should be verified as they may cause traffic to behave unexpectedly. In addition, decryption profiles should be checked to see if they are attached to a rule that is processing traffic, as this may lead to exceptions being made in handling. Also, various counters in the CLI can show CPU and memory utilization as it pertains to SSL decryption.

### 6.8.4 Authentication

System logs can show details of what authentication profile is used to authenticate a user. In the event of a failure, this information may also assist with identifying if there are problems with certificates, authentication methods, or communication.

### 6.8.5 References

Policies > NAT,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/policies/policies-nat>

Policies > Security,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/policies/policies-security>

Policies > Decryption,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/policies/policies-decryptio>

n

Policies > Authentication,

<https://docs.paloaltonetworks.com/pan-os/11-0/pan-os-web-interface-help/policies/policies-authentication>

## 6.9 Troubleshoot HA functions

### 6.9.1 Monitor

Link and path monitoring are used to verify if an HA peer is available. In the event of a failure, information will be written to the system log.

## 6.9.2 Failover triggers

When a failure occurs on one firewall and the peer in the HA pair (or a peer in the HA cluster) takes over the task of securing traffic, the event is called a *failover*. A failover is triggered, for example, when a monitored metric fails on a firewall in the HA pair. The metrics that the firewall monitors for detecting a firewall failure are:

- **Heartbeat Polling and Hello messages**

The firewalls use hello message and heartbeats to verify that the peer firewall is responsive and operational. Hello messages are sent from one peer to the other at the configured `Hello Interval` to verify the state of the firewall. The heartbeat is an ICMP ping to the HA peer over the control link, and the peer responds to the ping to establish that the firewalls are connected and responsive.

- **Link Monitoring**

You can specify a group of physical interfaces that the firewall will monitor (a link group) and the firewall monitors the state of each link in the group (link up or link down).

- **Path Monitoring**

You can specify a destination IP group of IP addresses that the firewall will monitor. The firewall monitors the full path through the network to mission-critical IP addresses using ICMP pings to verify reachability of the IP address.

In addition to the failover triggers listed above, a failover also occurs when the administrator suspends the firewall or when preemption occurs.

## 6.9.3 References

- HA Overview, <https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-admin/high-availability/ha-overview>
- Failover Triggers, <https://docs.paloaltonetworks.com/pan-os/10-2/pan-os-admin/high-availability/ha-concepts/failover>

## Appendix D: Glossary

- **Advanced Encryption Standard (AES):** A symmetric block cipher based on the Rijndael cipher.
- **Application Programming Interface (API):** A set of routines, protocols, and tools for building software applications and integrations.
- **Attack vector:** A path or tool that an attacker uses to target a network.
- **Boot sector:** Contains machine code that is loaded into an endpoint's memory by firmware during the startup process, before the operating system is loaded.
- **Boot sector virus:** Targets the boot sector or master boot record (MBR) of an endpoint's storage drive or other removable storage media. Also see boot sector and master boot record (MBR).

- **Bot:** Individual endpoints that are infected with advanced malware that enables an attacker to take control of the compromised endpoint. Also known as a zombie. See also botnet.
- **Botnet:** A network of bots (often tens of thousands or more) working together under the control of attackers using numerous command and control (C2) servers. See also bot.
- **Bring your own apps (BYOA):** Closely related to BYOD, BYOA is a policy trend in which organizations permit end users to download, install, and use their own personal apps on mobile devices, primarily smartphones and tablets, for work-related purposes. See also bring your own device (BYOD).
- **Bring your own device (BYOD):** A policy trend in which organizations permit end users to use their own personal devices, primarily smartphones and tablets, for work-related purposes. BYOD relieves organizations from the cost of providing equipment to employees but creates a management challenge due to the vast number and type of devices that must be supported. See also bring your own apps (BYOA).
- **Bulk electric system (BES):** The large interconnected electrical system, consisting of generation and transmission facilities (among others), that comprises the “power grid.”
- **Child process:** In multitasking operating systems, a sub-process created by a parent process that is currently running on the system.
- **Consumerization:** A computing trend that describes the process that occurs as end users increasingly find personal technology and apps that are more powerful or capable, more convenient, less expensive, quicker to install, and easier to use, than enterprise IT solutions.
- **Covered entity:** Defined by HIPAA as a healthcare provider that electronically transmits PHI (such as doctors, clinics, psychologists, dentists, chiropractors, nursing homes, and pharmacies), a health plan (such as a health insurance company, health maintenance organization, company health plan, or government program including Medicare, Medicaid, military and veterans’ healthcare), or a healthcare clearinghouse. See also Health Insurance Portability and Accountability Act (HIPAA) and protected health information (PHI).
- **Critical Infrastructure Protection (CIP):** Cybersecurity standards defined by NERC to protect the physical and cyber assets necessary to operate the bulk electric system (BES). See also the bulk electric system (BES) and North American Electric Reliability Corporation (NERC).
- **Data encapsulation:** A process in which protocol information from the OSI layer immediately above is wrapped in the data section of the OSI layer immediately below. See also open systems interconnection (OSI) reference model.
- **Distributed denial-of-service (DDOS):** A type of cyberattack in which extremely high volumes of network traffic such as packets, data, or transactions are sent to the target victim’s network to make their network and systems (such as an e-commerce website or other web application) unavailable or unusable.

- **Electronic health record (EHR):** As defined by HealthIT.gov, an EHR “goes beyond the data collected in the provider’s office and includes a more comprehensive patient history. EHR data can be created, managed, and consulted by authorized providers and staff from across more than one healthcare organization.”
- **Electronic medical record (EMR):** As defined by HealthIT.gov, an EMR “contains the standard medical and clinical data gathered in one provider’s office.”
- **Endpoint:** A computing device such as a desktop or laptop computer, handheld scanner, point-of-sale (POS) terminal, printer, satellite radio, security or video conferencing camera, self-service kiosk, server, smart meter, smart TV, smartphone, tablet, or Voice over Internet Protocol (VoIP) phone. Although endpoints can include servers and network equipment, the term is generally used to describe end user devices.
- **Enterprise 2.0:** A term introduced by Andrew McAfee and defined as “the use of emergent social software platforms within companies, or between companies and their partners or customers.” See also Web 2.0.
- **Exclusive OR (XOR):** A Boolean operator in which the output is true only when the inputs are different (for example, TRUE and TRUE equals FALSE, but TRUE and FALSE equals TRUE).
- **Exploit:** A small piece of software code, part of a malformed data file, or a sequence (string) of commands, that leverages a vulnerability in a system or software, causing unintended or unanticipated behavior in the system or software.
- **Extensible authentication protocol (EAP):** A widely used authentication framework that includes approximately 40 different authentication methods.
- **Extensible authentication protocol-Transport layer security (EAP-TLS):** An Internet Engineering Task Force (IETF) open standard that uses the Transport Layer Security (TLS) protocol in Wi-Fi networks and PPP connections. See also point-to-point protocol (PPP) and Transport Layer Security (TLS).
- **Extensible markup language (XML):** A programming language specification that defines a set of rules for encoding documents in a human- and machine-readable format.
- **False negative:** In anti-malware, malware that is incorrectly identified as a legitimate file or application. In intrusion detection, a threat that is incorrectly identified as legitimate traffic.
- **False positive:** In anti-malware, a legitimate file or application that is incorrectly identified as malware. In intrusion detection, legitimate traffic that is incorrectly identified as a threat.
- **Favicon (“favorite icon”):** A small file containing one or more small icons associated with a particular website or web page.
- **Federal Information Security Modernization Act (FISMA):** A U.S. law that implements a comprehensive framework to protect information systems used in U.S. federal government agencies. Known as the Federal Information Security Management Act prior to 2014. Financial Services Modernization Act of 1999: See Gramm-Leach-Bliley Act (GLBA).

- **Floppy disk:** A removable magnetic storage medium commonly used from the mid-1970s until approximately 2007, when they were largely replaced by removable USB storage devices.
- **Generic routing encapsulation (GRE):** A tunneling protocol developed by Cisco Systems® that can encapsulate various network layer protocols inside virtual point-to-point links.
- **Gramm-Leach-Bliley Act (GLBA):** A U.S. law that requires financial institutions to implement privacy and information security policies to safeguard the non-public personal information of clients and consumers. Also known as the Financial Services Modernization Act of 1999.
- **Hacker:** Originally used to refer to anyone with highly specialized computing skills, without connoting good or bad purposes. However, common misuse of the term has redefined a hacker as someone that circumvents computer security with malicious intent, such as a cybercriminal, cyberterrorist, or hacktivist.
- **Hash signature:** A cryptographic representation of an entire file or program's source code.
- **Health Insurance Portability and Accountability Act (HIPAA):** A U.S. law that defines data privacy and security requirements to protect individuals' medical records and other personal health information. See also covered entity and protected health information (PHI).
- **Heap spraying:** A technique used to facilitate arbitrary code execution by injecting a certain sequence of bytes into the memory of a target process.
- **Indicator of Compromise (IOC):** A network or operating system (OS) artifact that provides a high level of confidence that a computer security incident has occurred.
- **Initialization vector (IV):** A random number used only once in a session, in conjunction with an encryption key, to protect data confidentiality. Also known as a nonce.
- **Jailbreaking:** Hacking an Apple® iOS device to gain root-level access to the device. This is sometimes done by end users to allow them to download and install mobile apps without paying for them, from sources, other than the App Store®, that are not sanctioned and/or controlled by Apple®. Jailbreaking bypasses the security features of the device by replacing the firmware's operating system with a similar, albeit counterfeit version, which makes it vulnerable to malware and exploits. See also rooting.
- **Least privilege:** A network security principle in which only the permission or access rights necessary to perform an authorized task are granted.
- **Malware:** Malicious software or code that typically damages, takes control of, or collects information from an infected endpoint. Malware broadly includes viruses, worms, Trojan horses (including Remote Access Trojans, or RATs), anti-AV, logic bombs, backdoors, rootkits, bootkits, spyware, and (to a lesser extent) adware.
- **Master Boot Record (MBR):** Contains information about how the logical partitions (or file systems) are organized on the storage media, and an executable boot loader that starts up the installed operating system.



- **Metamorphism:** A programming technique used to alter malware code with every iteration, to avoid detection by signature-based anti-malware software. Although the malware payload changes with each iteration – for example, by using a different code structure or sequence, or inserting garbage code to change the file size – the fundamental behavior of the malware payload remains unchanged. Metamorphism uses more advanced techniques than polymorphism. See also polymorphism.
- **Microsoft Challenge-handshake authentication protocol (MS-CHAP):** A protocol used to authenticate Microsoft Windows-based workstation, using a challenge-response mechanism to authenticate PPTP connections without sending passwords.
- **Mutex:** A program object that allows multiple program threads to share the same resource, such as file access, but not simultaneously.
- **Network and Information Security (NIS) Directive:** A European Union (EU) directive that imposes network and information security requirements – to be enacted by national laws across the EU within two years of adoption in 2016 – for banks, energy companies, healthcare providers and digital service providers, among others.
- **nonce:** See initialization vector (IV).
- **North American Electric Reliability Corporation (NERC):** A not-for-profit international regulatory authority responsible for assuring the reliability of the bulk electric system (BES) in the continental U.S., Canada, and the northern portion of Baja California, Mexico. See also bulk electric system (BES) and Critical Infrastructure Protection (CIP).
- **Obfuscation:** A programming technique used to render code unreadable. It can be implemented using a simple substitution cipher, such as an exclusive or (XOR) operation, or more sophisticated encryption algorithms, such as the Advanced Encryption Standard (AES). See also Advanced Encryption Standard (AES), exclusive or (XOR), and packer.
- **One-way (hash) function:** A mathematical function that creates a unique representation (a hash value) of a larger set of data in a manner that is easy to compute in one direction (input to output), but not in the reverse direction (output to input). The hash function cannot recover the original text from the hash value. However, an attacker could attempt to guess what the original text was and see if it produces a matching hash value.
- **Open systems interconnection (OSI) reference model:** Defines standard protocols for communication and interoperability using a layered approach in which data is passed from the highest layer (application) downward through each layer to the lowest layer (physical), then transmitted across the network to its destination, then passed upward from the lowest layer to the highest layer. See also data encapsulation.
- **Packer:** A software tool that can be used to obfuscate code by compressing a malware program for delivery, then decompressing it in memory at run time. See also obfuscation.
- **Packet Capture (PCAP):** A traffic intercept of data packets that can be used for analysis.

- **Password Authentication protocol (PAP):** An authentication protocol used by PPP to validate users with an unencrypted password. See also point-to-point protocol (PPP).
- **Payment Card Industry Data Security Standards (PCI DSS):** A proprietary information security standard mandated and administered by the PCI Security Standards Council (SSC), and applicable to any organization that transmits, processes, or stores payment card (such as debit and credit cards) information. See also PCI Security Standards Council (SSC).
- **PCI Security Standards Council (SSC):** Comprised of Visa, MasterCard, American Express, Discover, and JCB, the SSC maintains, evolves, and promotes PCI DSS. See also Payment Card Industry Data Security Standards (PCI DSS).
- **Personal Information Protection and Electronic Documents Act (PIPEDA):** A Canadian privacy law that defines individual rights with respect to the privacy of their personal information, and governs how private sector organizations collect, use, and disclose personal information during business.
- **Personally Identifiable Information:** Defined by the U.S. National Institute of Standards and Technology (NIST) as “any information about an individual maintained by an agency, including (1) any information that can be used to distinguish or trace an individual’s identity... and (2) any other information that is linked or linkable to an individual...”
- **Point-to-point protocol (PPP):** A Layer 2 (data link) protocol layer used to establish a direct connection between two nodes.
- **Polymorphism:** A programming technique used to alter a part of malware code with every iteration, to avoid detection by signature-based anti-malware software. For example, an encryption key or decryption routine may change with every iteration, but the malware payload remains unchanged. See also metamorphism.
- **Pre-shared key (PSK):** A shared secret, used in symmetric key cryptography which has been exchanged between two parties communicating over an encrypted channel.
- **Promiscuous mode:** Refers to Ethernet hardware used in computer networking, typically a network interface card (NIC), that receives all traffic on a network segment, even if the traffic is not addressed to the hardware.
- **Protected health information (PHI):** Defined by HIPAA as information about an individual’s health status, provision of healthcare, or payment for healthcare that includes identifiers such as names, geographic identifiers (smaller than a state), dates, phone and fax numbers, email addresses, Social Security numbers, medical record numbers, or photographs, among others. See also Health Insurance Portability and Accountability Act (HIPAA).
- **Public key infrastructure (PKI):** A set of roles, policies, and procedures needed to create, manage, distribute, use, store, and revoke digital certificates and manage public key encryption.
- **Quality of Service (QoS):** The overall performance of specific applications or services on a network including error rate, bit rate, throughput, transmission delay, availability, jitter, etc.

QoS policies can be configured on certain network and security devices to prioritize certain traffic, such as voice or video, over other, less performance-intensive traffic, such as file transfers.

- **Rainbow Table:** A pre-computed table used to find the original value of a cryptographic hash function.
- **Remote Authentication Dial-In User Service (RADIUS):** A client/server protocol and software that enables remote access servers to communicate with a central server to authenticate users and authorize access to a system or service.
- **Remote procedure call (RPC):** An inter-process communication (IPC) protocol that enables an application to be run on a different computer or network, rather than the local computer on which it is installed.
- **Representational state transfer (REST):** An architectural programming style that typically runs over HTTP, and is commonly used for mobile apps, social networking websites, and mashup tools.
- **Rooting:** The Google Android equivalent of jailbreaking. See jailbreaking.
- **Salt:** Randomly generated data that is used as an additional input to a one-way function that hashes a password or passphrase. The same original text hashed with different salts results in different hash values.
- **Sarbanes-Oxley (SOX) Act:** A U.S. law that increases financial governance and accountability in publicly traded companies.
- **Script kiddie:** Someone with limited hacking and/or programming skills that uses malicious programs (malware) written by others to attack a computer or network.
- **Secure Sockets Layer (SSL):** A cryptographic protocol for managing authentication and encrypted communication between a client and server to protect the confidentiality and integrity of data exchanged in the session.
- **Service set identifier (SSID):** A case sensitive, 32-character alphanumeric identifier that uniquely identifies a Wi-Fi network.
- **Software as a Service (SaaS):** A cloud computing service model, defined by the U.S. National Institute of Standards and Technology (NIST), in which “the capability provided to the consumer is to use the provider’s applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser, or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, except for limited user-specific application configuration settings.”
- **Spear phishing:** A highly targeted phishing attack that uses specific information about the target to make the phishing attempt appear legitimate.

- **Structured threat information expression (STIX):** An XML format for conveying data about cybersecurity threats in a standardized format. See also extensible markup language (XML).
- **Threat Vector:** See attack vector.
- **TLS:** See Transport Layer Security (TLS).
- **Tor (“The Onion Router”):** Software that enables anonymous communication over the internet.
- **Transport Layer Security (TLS):** The successor to SSL (although it is still commonly referred to as SSL). See also Secure Sockets Layer (SSL).
- **Uniform resource locator (URL):** A unique reference (or address) to an Internet resource, such as a web page.
- **Vulnerability:** A bug or flaw that exists in a system or software and creates a security risk.
- **Web 2.0:** A term popularized by Tim O’Reilly and Dale Dougherty, unofficially referring to a new era of the World Wide Web, which is characterized by dynamic or user-generated content, interaction, and collaboration, and the growth of social media. See also Enterprise 2.0.
- **Zero-day threat:** The window of vulnerability that exists from the time a new (unknown) threat is released until security vendors release a signature file or security patch for the threat.
- **Zombie:** See bot.

## Continuing Your Learning Journey with Palo Alto Networks

Training from Palo Alto Networks and our Authorized Training Partners delivers the knowledge and expertise to prepare you to protect our way of life in the digital age. Our trusted security certifications give you the Palo Alto Networks product portfolio knowledge necessary to prevent successful cyberattacks and to safely enable applications.

### Digital Learning

For those of you who want to keep up to date on our technology, a learning library of *free* digital learning is available. These on-demand, self-paced digital-learning classes are a helpful way to reinforce the key information for those who have been to the formal hands-on classes. They also serve as a useful overview and introduction to working with our technology for those unable to attend a hands-on, instructor-led class.

Simply register in [Beacon](#) and you will be given access to our digital-learning portfolio. These online classes cover foundational material and contain narrated slides, knowledge checks, and, where applicable, demos for you to access.

New courses are being added often, so check back to see new curriculum available.

## Instructor-Led Training

Looking for a hands-on, instructor-led course in your area?

Palo Alto Networks Authorized Training Partners (ATPs) are located globally and offer a breadth of solutions from onsite training to public, open-environment classes. About 42 authorized training centers are delivering online courses in 14 languages and at convenient times for most major markets worldwide. For class schedule, location, and training offerings, see <https://www.paloaltonetworks.com/services/education/atc-locations>.

## Learning Through the Community

You also can learn from peers and other experts in the field. Check out our communities site at <https://live.paloaltonetworks.com>, where you can:

- Discover reference material
- Learn best practices
- Learn what is trending



3000 Tannery Way  
Santa Clara, CA 95054

Main: +1.408.753.4000

Sales: +1.866.320.4788

Support: +1.866.898.9087

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

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