

Application DDoS Mitigation

©2014, Palo Alto Networks, Inc.

Contents

Overview 3

Volumetric vs. Application Denial of Service Attacks 3

Volumetric DoS Mitigation 4

Application Attack Mitigation 5

Signature Design Considerations 6
Custom Vulnerability Signature Creation Overview 6
Custom Signature Context 6

Application DoS Mitigation Examples 7

Excessive File Download Attack 7
Brute Force Authentication Attack 11
SSL Connection Exhaustion Attack 15
Excessive Search Attack 18
DNS Query Attack 20
Mitigation with Geo Locations 23

Summary 27

Revision History 27

©2014, Palo Alto Networks, Inc.

Overview

Distributed Denial of Service (DDoS) and Denial of Service (DoS) attacks can be classified in several categories and can be either volumetric or application based. By far, volumetric attacks have been the most popular types of DoS attacks, as readily available tools such as LOIC and HOIC can make it fairly easy to mount an attack. Application attacks that leverage *allowed* security policies have been gaining popularity as they are very difficult to guard against and can have the same detrimental effect in bringing down critical services.

This document will discuss the primary differences between volumetric and application DoS attacks and will show you how Palo Alto Networks custom vulnerability signatures and security policies can be used to effectively identify and mitigate application based DoS attacks.

Volumetric vs. Application Denial of Service Attacks

Volumetric DDoS/DoS attacks can take many forms, but the goals are the same. The attack attempts to take down services by flooding the victim service/device with a high volume of attack traffic which fully consumes all available resources, bandwidth, or both. Popular volumetric attack types include:

TCP Stack Flood	Attack that flood specific functions of the TCP connection process to keep a host from being able to respond to legitimate connections. Examples include TCP Syn, TCP Fin, TCP RST, TCP Flag attacks.
Generic Flood	Attack that floods one or more protocols and/or ports and are designed to resemble legitimate traffic and source addresses are often spoofed. Examples include: Ping Flood, Smurf, Reflection Attack, UDP Flood attacks.
Fragmentation Attack	Attack that sends a high volume of TCP or UDP fragments to a victim host. Designed to overwhelm the host's ability to re-assemble the packets and degrade its performance. Fragments can often be malformed to cause additional processing.
TCP Connection Attack	Attack that maintains a large number of half-open or full TCP connections to take up server resources and prevent legitimate sessions from completing.

Application DDoS/DoS attacks can also take many forms and they are generally targeted at a specific application or function. The goal is to disrupt services by using a legitimate application to create the attack. With application attacks, it's often very difficult for companies to block the malicious activity as the victim application is a legitimate function that is required to conduct business. The security and firewall policies are configured to allow access to the application being attacked. Popular types of application-based attacks include:

Application Attack	Attacks that are designed to overwhelm an application or specific components of an application. Common attacks are against HTTP, DNS, and SIP services. Often stealthy in nature and uses legitimate applications and ports to cause resource drain. The target may be an authentication server, a file download site, or some other critical service. Examples include HTTP GET Flood, SIP Invite Flood, DNS Amplification Attack, HTTP/FTP File Download, and SSL Connection attacks.
Vulnerability Exploits	Attack designed to take advantage of vulnerabilities in the operating system and/or application to gain unauthorized access. In many cases, the firewall's vulnerability policies can be used to spot and mitigate against these types of attacks.
SSL Attack	A form of application attack that is designed to leverage resource intensive SSL key negotiation or re-keying activities to drain the targeted victim's resources and bring down legitimate SSL services.

©2014, Palo Alto Networks, Inc.

Volumetric DoS Mitigation

Volumetric attacks are the most popular DDoS attacks as they are fairly easy to mount against a host, service or web site. They are also fairly easy to spot as they tend to generate traffic patterns outside normal usage patterns. In order to be successful with a volumetric attack, the attacker simply needs to accomplish one of the following:

1. Generate enough attack traffic to consume all processing, memory, or buffer resources on the victim host/service.
2. Consume enough bandwidth to fully congest the network connection to the host/service.
3. Generate enough attack traffic to consume all resources of the network/security equipment protecting the victim host/service.

In general, volumetric protection strategies should involve a layered approach and be comprised of local DDoS/DoS protection at the enterprise/data center location and protection at the service provider and/or cloud location. The two pronged strategy protects the Internet connection (bandwidth) with ISP/cloud policies and specific services with local firewall DDoS/DoS policies. For example, if only local DoS protection is implemented in front of the data center, the attacker can simply launch a volumetric attack that consumes all of the bandwidth connecting the service to the Internet. Regardless if the local DoS solution can stop the attack or not, the attacker has achieved their goal as the service is down due to the lack of bandwidth to the Internet.

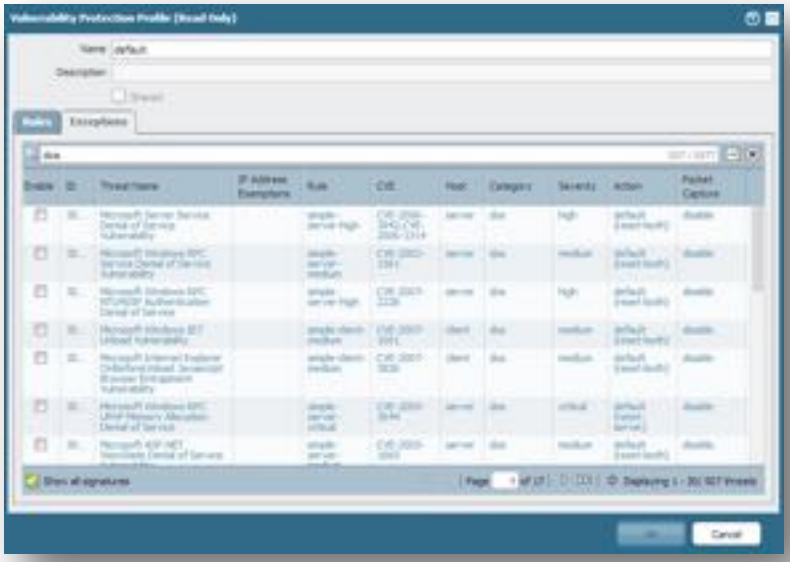
Note: Some things to investigate with cloud-based DDoS solutions are the additional latency and/or bandwidth that may be experienced when a volumetric attack is underway. For example, some cloud solutions require the redirection of affected traffic from your site to their cloud *traffic scrubbers* through DNS redirection or GRE tunnels.

As part of successful denial of service mitigation strategy, the following PAN-OS features can be implemented locally to help secure critical resources.

DoS Profiles and Polices	Mitigates against SYN, UDP, ICMP, ICMPv6, and Other IP Flood attacks. Provides resource protection by limiting the number of sessions that can be used. Implemented on a per-source bases (classified) or on an aggregate bases and can be applied to specific, zones, interfaces, IP address(es), user(s), and service(s). Source IP addresses can be geo-location based to allow flexible DoS policies that limit specific countries more than others. All DoS policies can also be scheduled.
Zone Profiles	Mitigates against SYN, UDP, ICMP, ICMPv6, and Other IP Flood attacks on a zone bases, regardless of interface(s). Provides reconnaissance protection against TCP/UDP port scans and host sweeps as well as packet-based attacks that leverage TCP header options, spoofed IP addresses, fragmentation, mismatched overlapping TCP segments, ICMP/ICMPv6 options, and IPv6 options. Zone protection is applied to a Zone versus specific interfaces or IP addresses to provide a *broad first pass* level of protection against malformed packets and common flood attacks.
Vulnerability Protection Profiles	Volumetric attacks that leverage known vulnerabilities, brute force methods, amplification, spoofing, and others can be effectively blocked through PAN-OS vulnerability protection profiles. To obtain a list of the latest Vulnerability Protection DoS capabilities, you can use one of two methods: 1) Visit the Palo Alto Networks Threat Vault and search the threat database with the keyword *DoS*. Use the *Type* dropdown box to view vulnerability, spyware, and virus related DoS attacks.



2) On the firewall web UI's *Objects* tab, open the *Default* vulnerability protection profile. Under the *Exceptions* tab, check the *Show all signatures* option and filter on the keyword *DoS* as shown in the following illustration.



Note: Reference the PAN-OS Administrators Guide for more information on Zone and DoS mitigation features.

Application Attack Mitigation

Application DDoS/DoS are attacks that use approved applications and protocols to victimize a specific host or service. The attack's goal is to consume enough resources to eventually disrupt legitimate traffic and the attacker relies on the security policy's *permit* action to gain access. This makes application attacks much harder to defend against than volumetric attacks. What makes matters worse, is the ability to accurately differentiate attack traffic from legitimate traffic and to identify which sessions belong to the attacker and which do not.

In these situations, there is a difficult choice for security teams.

- Do they implement security policies to keep the service up at all costs and sacrifice a percentage of legitimate user traffic when the attack is blocked?
- Or do they design the network and services with enough capacity to absorb the attack as blocking legitimate users (false positives) are not permitted?

If the strategy is to address the first concern, Palo Alto Networks firewalls can be configured with custom vulnerability signatures to identify suspicious application activity and proactively mitigate DDoS/DoS attacks. Based on integer context, string context, and regex comparisons, combination signatures are a powerful tool to enforce specific thresholds to proactively protect against resource depletion. Context filters can include time attributes to monitor services for abnormal activity – such as the number of HTTP request containing the word "password" exceeding 50 transactions per second.

Signature Design Considerations

When designing custom signatures for application attack policies, several things to keep in mind include:

1. Aggregation criteria should be considered carefully. Many users (source IP address) may originate sessions behind network translated (NAT) firewalls and hence, the traffic of many users may appear to be coming from the same source IP address. This should be taken into consideration when defining the thresholds.
2. When protecting a server/service, the capacity of that service should be known before creating the combination signature's threshold. For example, if the FTP service can handle an average of 50 connections per second before maxing out, the connection threshold should be set lower to prevent resource exhaustion.
3. When the specified threshold is reached, the signatures default action will activate (alert, drop packets, reset client, reset server, reset both). Remember that both legitimate and attack traffic will be impacted when protection is activated.

Custom Vulnerability Signature Creation Overview

The application DDoS/DoS mitigation examples in this document will use custom vulnerability signatures to identify the application/service and will use *combination signatures* to add thresholds to monitor and enforce activity. The process to create and activate custom vulnerability signatures in security policies is the same between examples and includes the following steps:

1. Create the custom vulnerability signature to identify the application/service.
2. Create a custom combination signature to set the time and aggregation criteria to define how thresholds are applied to the application/service.
3. Enable *Vulnerability Protection* on the security policy that controls access to the application/service.

By default, all custom vulnerability signatures created are automatically enabled when vulnerability protection profiles are enabled on security policies. You can optionally create vulnerability protection profiles to exclude any signature that should not be enforced.

Custom Signature Context

Custom signatures use context to match traffic. PAN-OS provides the following context filter capabilities and protocol support.

- Integer Context – Greater Than, Less Than, Equal To
  - FTP
  - HTTP
  - IMAP
  - SMTP
- String Contexts - Pattern Match
  - DNS
  - Email, IMAP, SMTP
  - File Body
  - FTP
  - GDB Debugger
  - GIOP
  - HTTP
  - IRC
  - JPEG
  - MS SMB, RPC, SQL
  - PE Header
  - RTMP, RTSP
  - SSH, SSL
  - Telnet
  - Unknown
- Regex Syntax

Tip: Reference the *Creating Custom Signatures Tech Note – PAN-OS 5.0* document for more information on creating custom vulnerability signatures.

Application DoS Mitigation Examples

The use cases illustrated in this paper are actual attacks customers have experienced. The mitigation techniques are designed to show the flexibility and power of custom signatures and how they can be applied on the firewall to proactively identify and defend against possible attacks. These use cases should be used as a reference to help model mitigation strategies within your environment to help secure critical resources.

Excessive File Download Attack

A file download attack can be described as a server resource exhaustion and/or bandwidth consumption attack. Before launching the attack, reconnaissance of the target service will be performed by the attacker to discover the files that can be downloaded along with the file sizes. For example, they may probe a company's download site for their largest PDF file. Using this information, the attacker will leverage hundreds, or even thousands, of bots to execute the attack and attempt to download the large file simultaneously. As a file download from the company's web site is a valid operation, the security policies will allow this activity. The attack succeeds if it can either overwhelm the download server, consume all available bandwidth, or both.

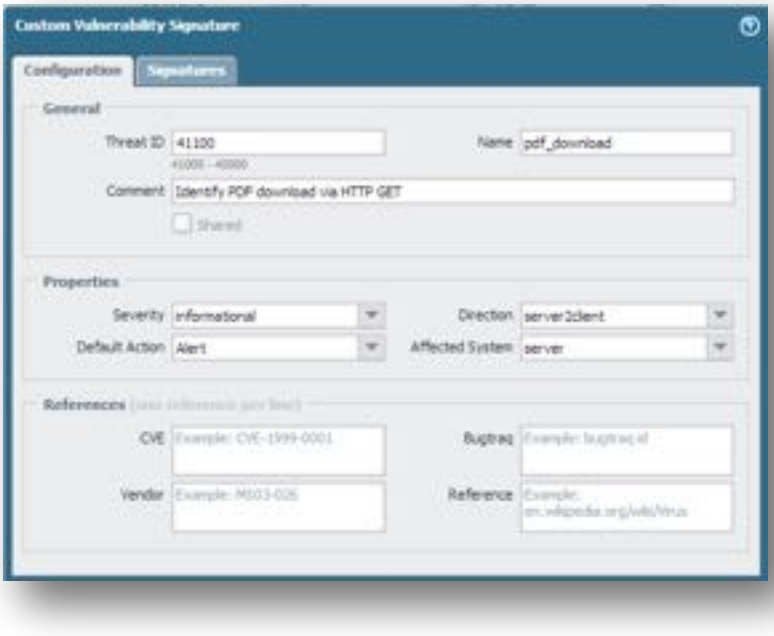
To monitor and mitigate against this type of attack, a proactive custom vulnerability signature can be created to guard the download server against high levels of simultaneous file downloads. For example, if an HTTP GET was used to download the file(s), the signature will monitor the number of HTTP GET transactions issued against the download server's IP address and download folder. A threshold will define the defensive action taken when the volume of downloads are exceeded.

In this example, a custom vulnerability signature will be created to identify the download activity and it will be used within a combination signature to enforce the permitted capacity of 35 downloads per second.

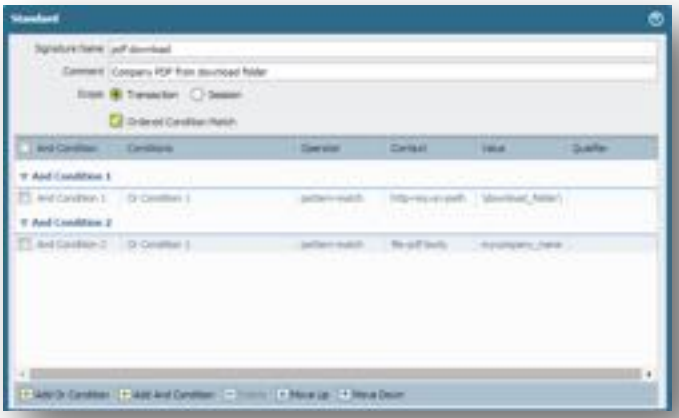
Deployment Example:

1. Create a custom vulnerability signature using the HTTP GET context to identify PDF downloads from a specific folder. This example shows custom signature with Threat ID 41100 and the following characteristics.

    a. Configuration Tab defines the traffic properties:
    Severity = informational
    Direction = server2client
    Default Action = Alert
    Affected System = server



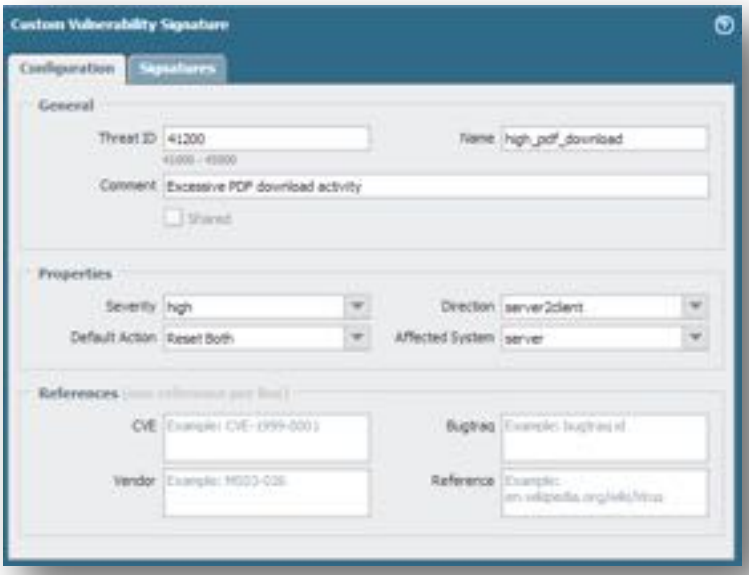
b. Signatures Tab defines the match criteria:
   Scope = Transaction
   Context = http-req-uri-path (pattern match) for download_folder
   Context = file-pdf-body (pattern match) for ABC Corporation



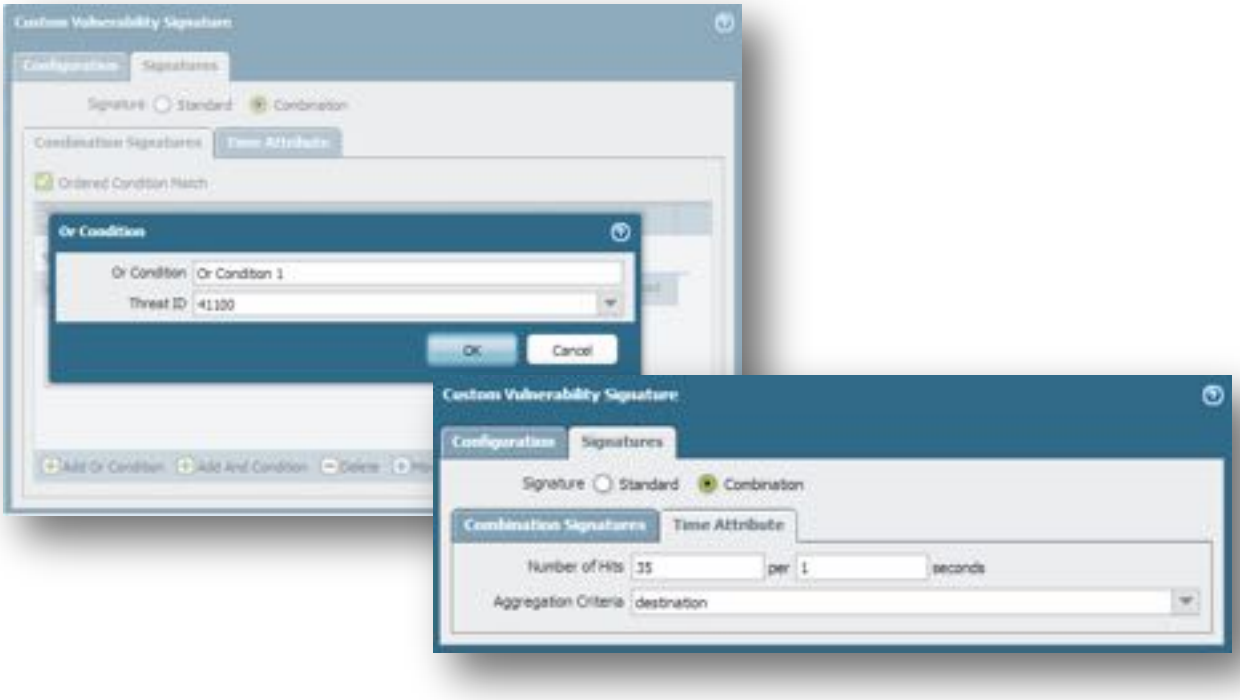
To summarize, this custom signature will look for any PDF file transactions containing *ABC Corporation* originating from the server's *download_folder* in the URI string. The *Transaction* scope setting allows the signature to match on individual transactions versus sessions to spot multiple downloads within the same TCP session. When there is a match, a log event is created in the Threat Log.

2. Create a custom combination signature to combine the custom signature from the previous step and to set the time attributes for the protection threshold. This example shows custom signature with Threat ID 41200 and the following characteristics:

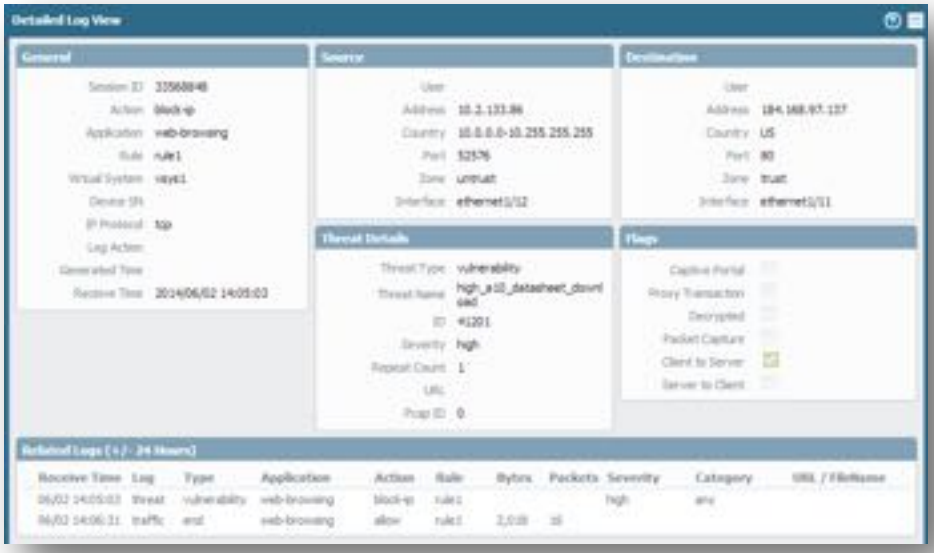
   a. Configuration Tab defines the traffic & action properties:
      Severity = high
      Direction = server2client
      Default Action = Reset Both (resets client and server)
      Affected System = server



   b. Signatures Tab defines the protection threshold:
      Signature = Combination
      Threat ID = specifies the custom vulnerability signature from step 1 (41100)
      Time Attribute = 35 (set PDF download threshold to 35 times per 1 second interval)
      Aggregation Criteria = destination (from server side)

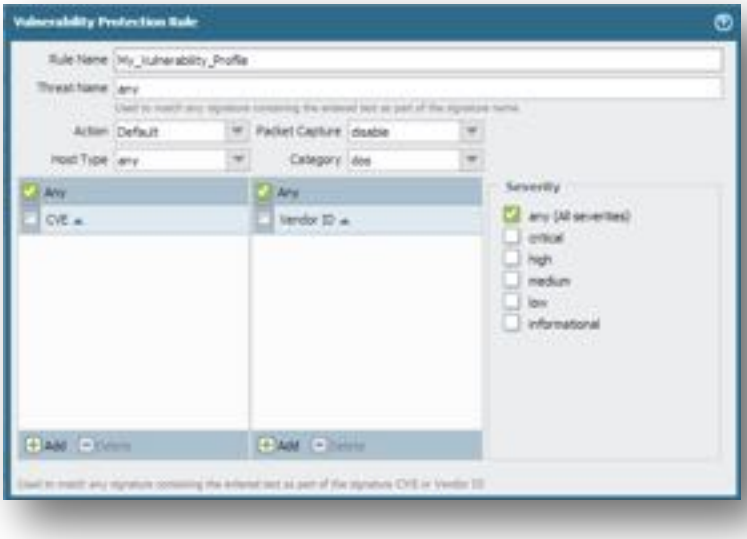


To summarize, this combination signature will trigger when corporate PDFs from the *download_folder* is seen more than 35 times per second originating from the server side. When the threshold is reached, the action will be to reset both the client and server to disable the download attempt. An event with a *High* severity is logged in the Threat Log.

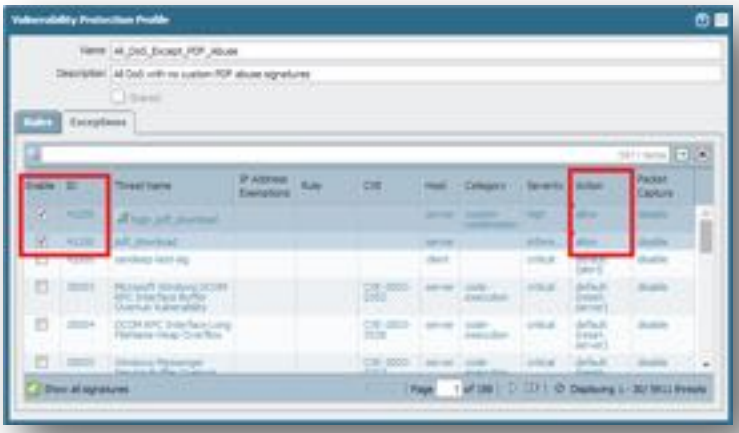


3. By default, custom vulnerability signatures that you create are automatically enabled when they are added to the vulnerability database. Optionally, you can create a vulnerability protection profile to exclude the custom signatures created in steps 1 and 2 and apply it to security policies that do not need them. To exclude a custom vulnerability signature using a new Vulnerability Protection Profile, perform the following:

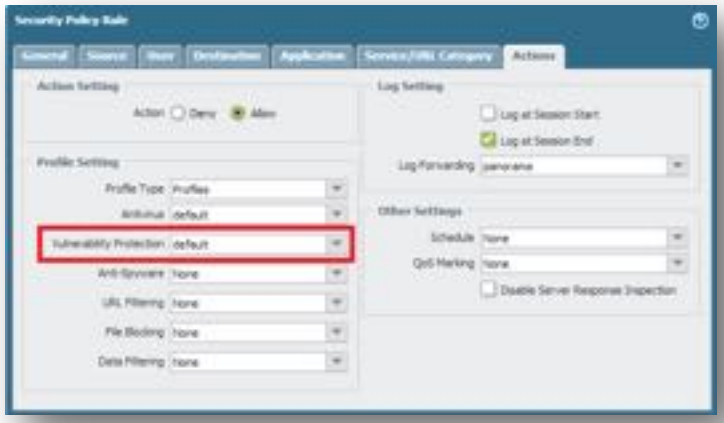
    a. Create a new rule with the desired vulnerability filters. In this example, we will create a rule to include all *DoS* category vulnerabilities with *Any* severity and set the action to Default to accept the Palo Alto Networks preconfigured actions.



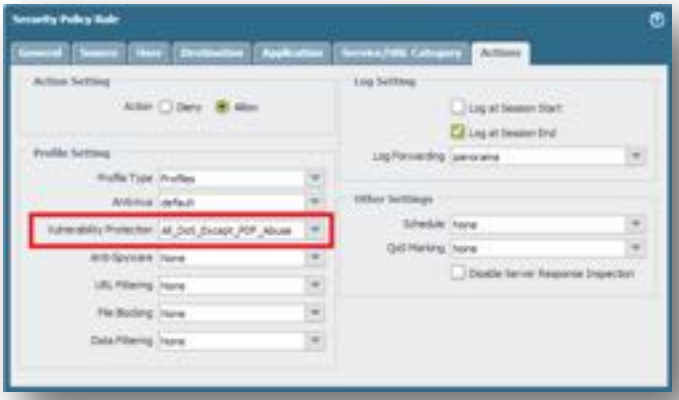
b. As PDF download attack detection and mitigation is not required in this protection profile, exclude the combination signatures created in the previous two steps. Under the *Exceptions* tab, check the *Show all signatures* checkbox to display all vulnerability signatures. Locate the two custom signatures (41100 and 41200) and enable them for exception with the left most checkbox under the *Enable* column. Under the *Action* column, set the action to *allow* to bypass these two signatures.



4. In order to activate the custom signatures, enable *Vulnerability Protection* on the security policy controlling access to the download application/servers. All custom signatures are automatically enabled with the pre-defined (default and strict) vulnerability protection profiles unless they are disabled using the exceptions function.



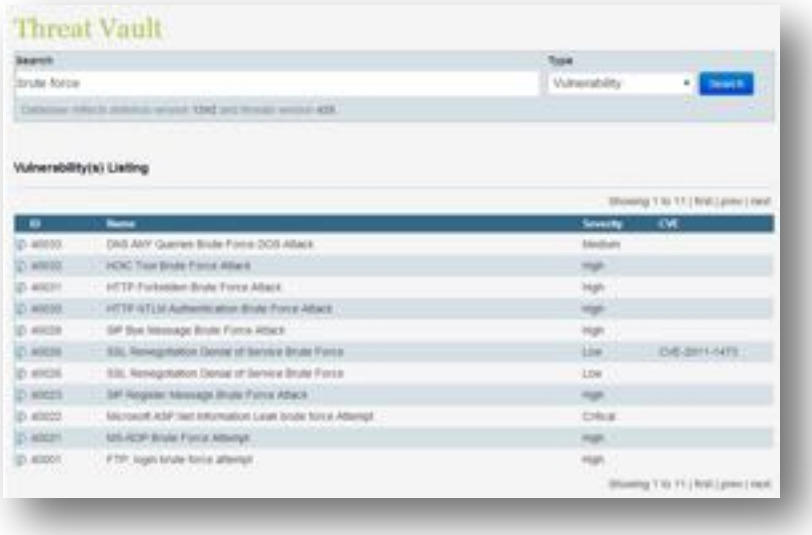
5. If you need to bypass the custom vulnerability signatures, select the vulnerability protection profile that was created in step 3 and apply it to the security policies that do not need the custom signatures.



Tip: Dynamic Address Groups can be used in security policies to simplify management of server objects and eliminate the need to change IP addresses when servers move or change - such as in a hypervisor environment where new servers are dynamically added and removed.

Brute Force Authentication Attack

Well-known brute force attacks can be prevented with PAN-OS's pre-defined vulnerability signatures and profiles. To view the latest Brute Force attack signatures, search the Palo Alto Networks Threat Vault with the *brute force* search criteria as shown in the following illustration.

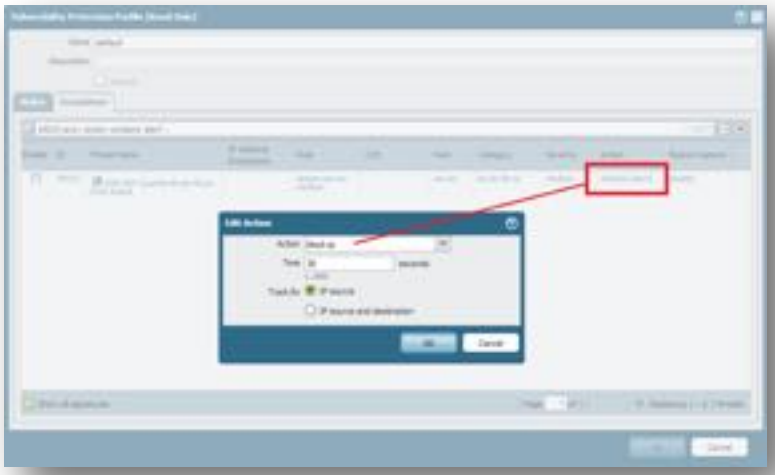


Example of Brute Force Protection through pre-defined PAN-OS Vulnerability Signatures

Modifying Pre-defined Vulnerability Signatures

If you need to modify any of the pre-defined vulnerability signature settings, such as the action to take when the threat is detected, a new Vulnerability Protection Profile with an exception can be created and applied to the security policy. The following example illustrates how to change the action from *default (alert)* to *block-ip* for a 30 second duration for a predefined signature.

1. Create a new vulnerability profile by cloning one of the firewall's pre-defined profiles (default or strict profile) or clicking on the Add button.
2. On the *Exceptions* tab, select the *Show all signatures* checkbox to display all vulnerability signatures. You can also use the search function if you know the signature name or CVE ID.
3. Select the signature with the *Enable* checkbox and modify the *Action* with the dropdown list. You can select from the various action types to change the default behavior.
4. Save the changes and apply the vulnerability protection profile to the security policy.



Creating Custom Brute Force Signatures

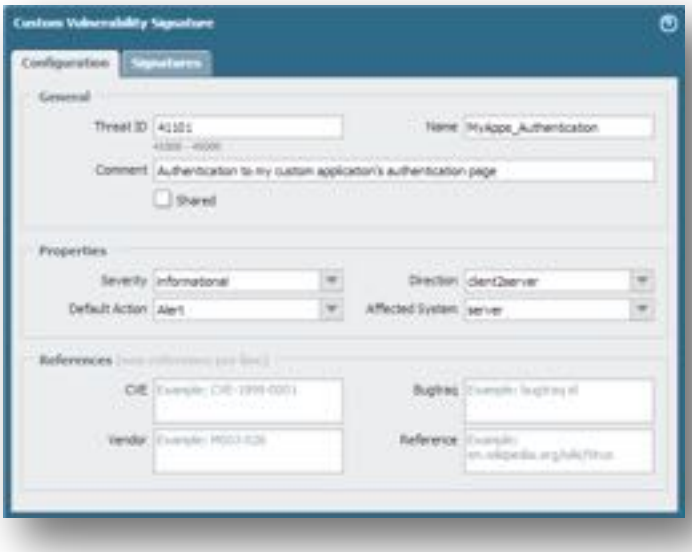
If the pre-defined brute force vulnerability signatures do not meet your needs, custom signatures can be created to help detect specific brute force attacks. The following example illustrates how a custom vulnerability signature is used to protect an application's authentication page from a dictionary attack by monitoring the number of HTTP POSTs originating from the source IP address.

Tip: Use the PAN-OS decryption feature to analyze encrypted traffic.

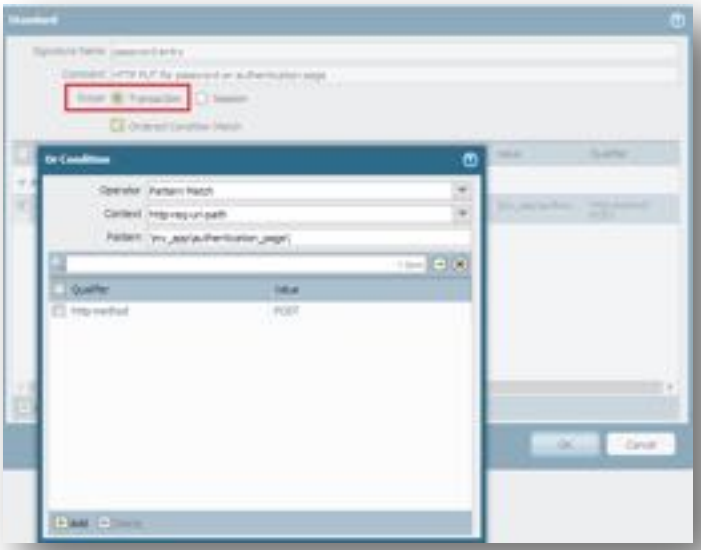
Deployment Example:

1. Create a custom vulnerability signature for the application's authentication function by matching on the authentication page's URI string and the HTTP PUT qualifier for the user credentials. This example shows custom signature with Threat ID 41101 and the following characteristics.

    a. Configuration Tab defines the traffic properties:
        Severity = informational
        Direction = client2server
        Default Action = Alert
        Affected System = server



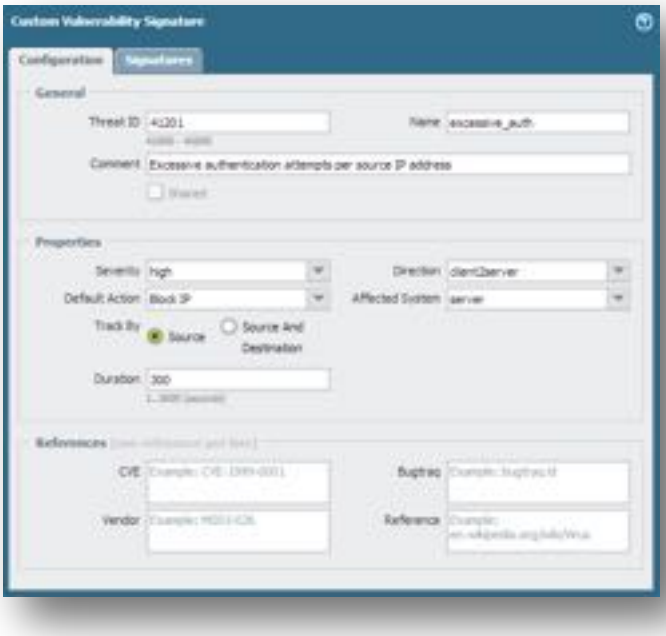
b. Signatures Tab defines the match criteria:
   Scope = Transaction
   Operator = Pattern Match
   Context = http-req-uri-path
   Pattern = my_app\authentication_page
   Qualifier (http-method) = POST



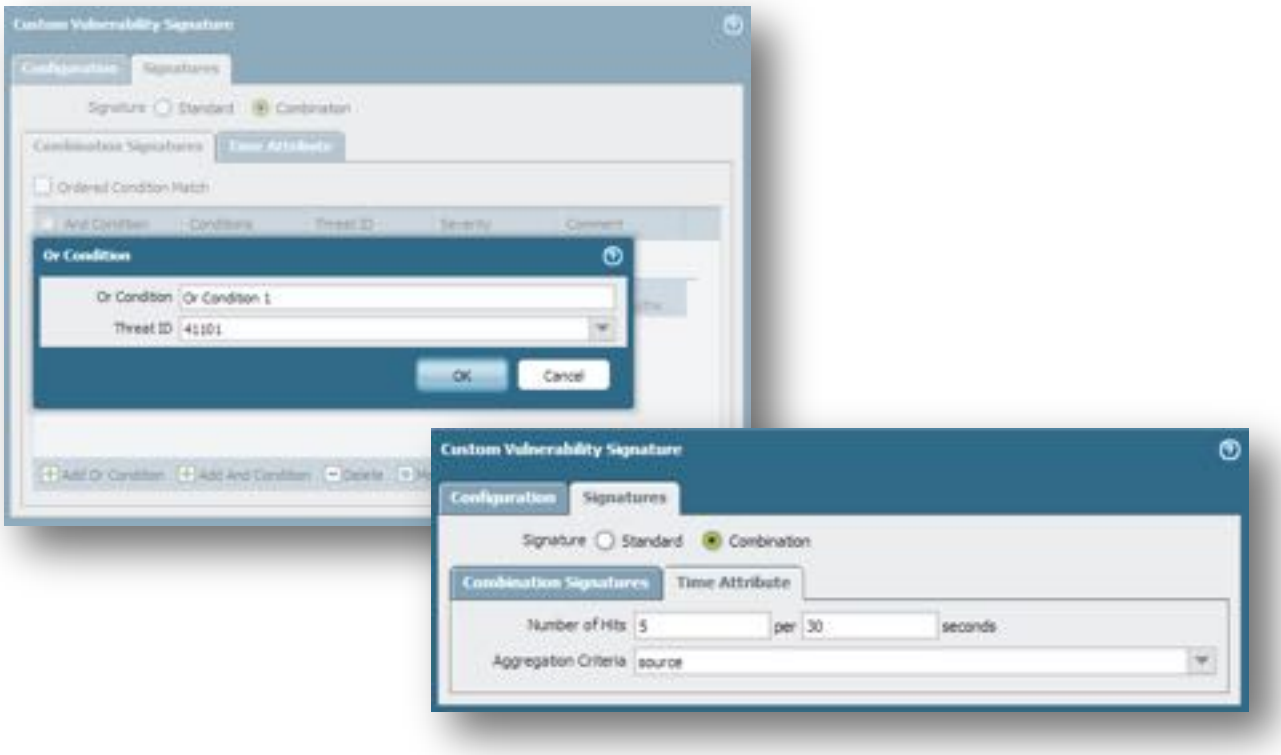
To summarize, this custom signature will look for HTTP POST transactions in the URI path containing *my_app\authentication_page*. The *Transaction* scope setting allows the signature to match on individual transactions versus sessions to spot multiple authentication attempts within the same TCP session. When there is a match, a log event is created in the Threat Log with the *Alert* action and *informational* severity.

2. Create a custom vulnerability *combination signature* to combine the custom signature from the previous step and to set the time attributes for the protection threshold. This example shows custom signature with Threat ID 41201 and the following characteristics:

   a. Configuration Tab defines the traffic & action properties:
      Severity = high
      Direction = client2server
      Default Action = Block IP
      Affected System = server
      Track By = Source
      Duration = 300 seconds (block duration)

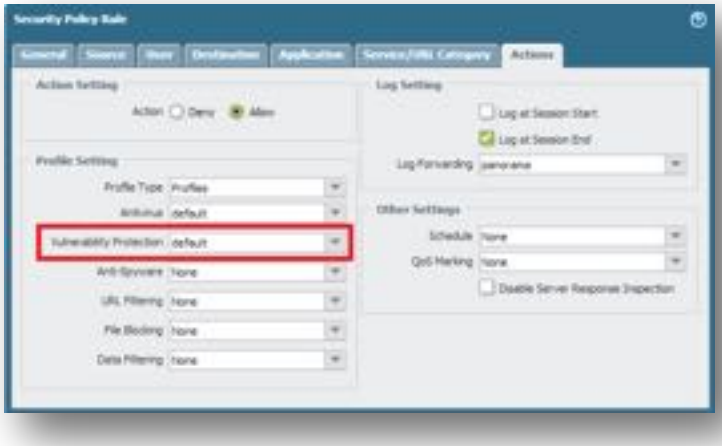


   b. Signatures Tab defines the protection threshold:
      Signature = Combination
      Threat ID = specifies the custom vulnerability signature from step 1 (41101)
      Time Attribute = 5 (sets authentication threshold to 5 times per 30 second interval)
      Aggregation Criteria = source (from client side)



To summarize, this combination signature will trigger when the same source IP address performs more than five HTTP PUT transactions on the authentication page within 30 seconds. When the threshold is reached, the action will be to block the offending source IP address for 300 seconds. An event with a *High* severity is logged in the Threat Log.

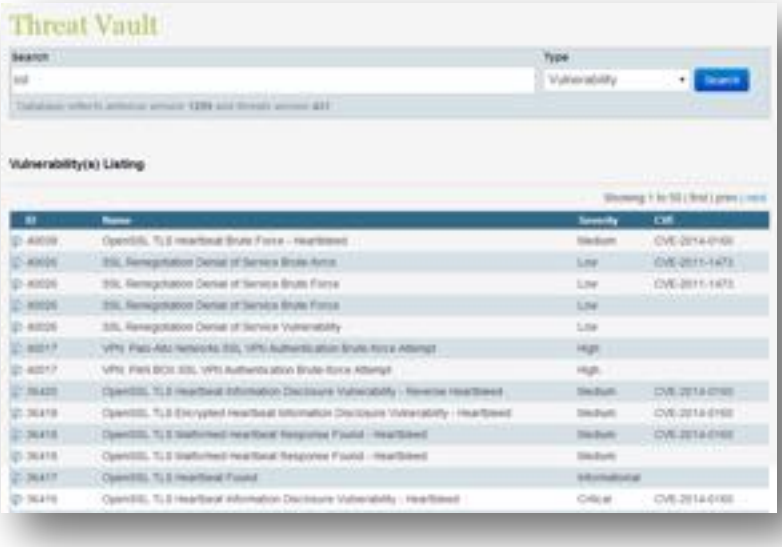
3. In order to activate the custom signatures, enable the *Vulnerability Protection* option on the security policy controlling access to the authentication page. All custom signatures are automatically enabled with the pre-defined (default and strict) vulnerability protection profiles.



4. If you need to bypass the custom vulnerability signatures for specific security policies, you can create a new vulnerability protection profile and exclude the custom signatures. See steps 3 and 5 from the first example for more details.

SSL Connection Exhaustion Attack

PAN-OS's pre-defined vulnerability signatures provide a wide range of protection against SSL attacks (such as SSL Renegotiation Brute Force, SSL VPN Authentication Brute Force, and others). However, there may be cases where these signatures don't provide the exact detection and mitigation for an SSL application. To view the latest SSL vulnerability signatures, search the Palo Alto Networks Threat Vault with the *SSL* search string as shown in the following illustration.



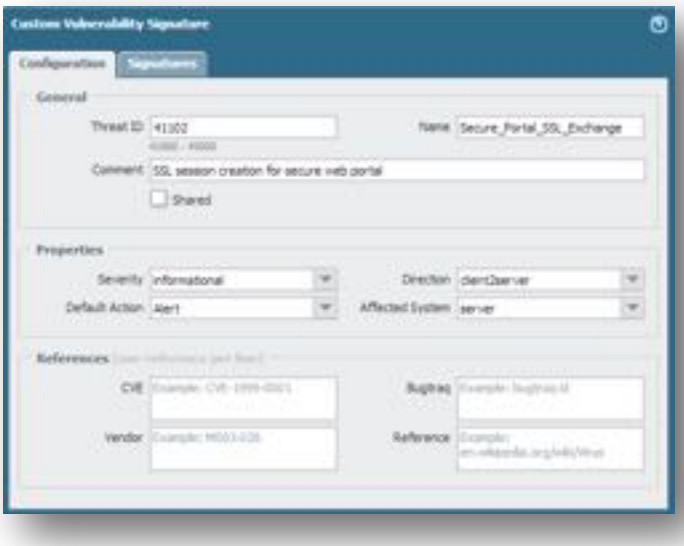
One of the easier SSL application attacks that can be launched against an SSL server/service is a resource exhaustion attack. Through the use of a botnet, thousands of new SSL connections can be made to the SSL server over a predefined time interval. Each new SSL connection requires server resources to generate keys and encrypt the session and the attack simply increases the number of SSL connections until the desired server latency is reached.

The following example illustrates how a custom vulnerability signature can be used to enforce the number of SSL connection attempts per source IP address. This can help reduce the chances of an SSL attack that is designed to exhaust server CPU resources with empty SSL sessions.

Deployment Example:

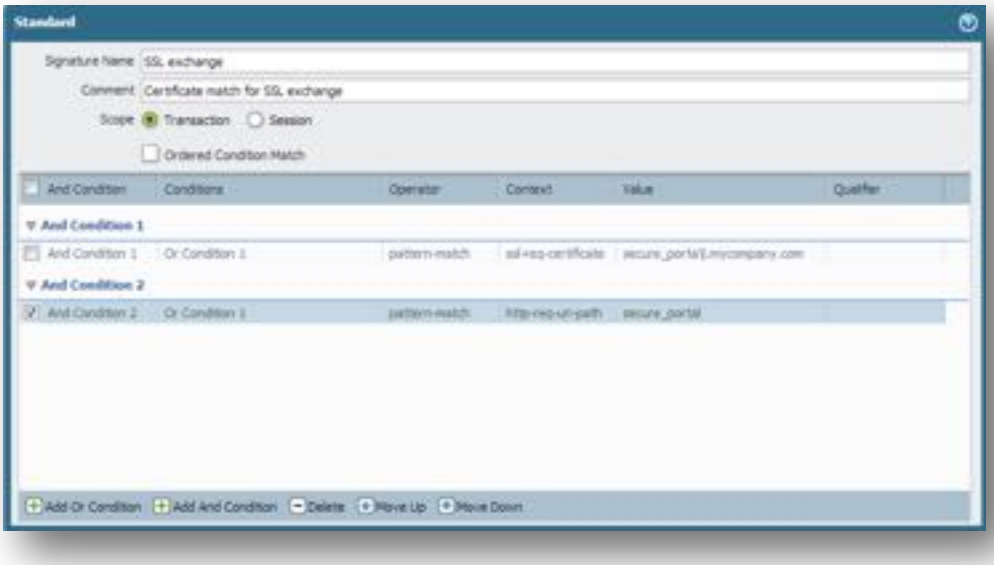
1. Create a custom vulnerability signature for the SSL application by matching on the FQDN field of the certificate used in the SSL negotiation exchange and the portal's URI string. This example shows custom signature with Threat ID 41102 and the following characteristics.

   a. Configuration Tab defines the traffic properties:
      Severity = informational
      Direction = client2server
      Default Action = Alert
      Affected System = server



b. Signatures Tab defines the match criteria for two items:
   Scope = Transaction
   Operator = pattern-match
   Context 1 = ssl-req-certificate
   Value = secure_portal.mycompany.com (FQDN identifier in certificate)

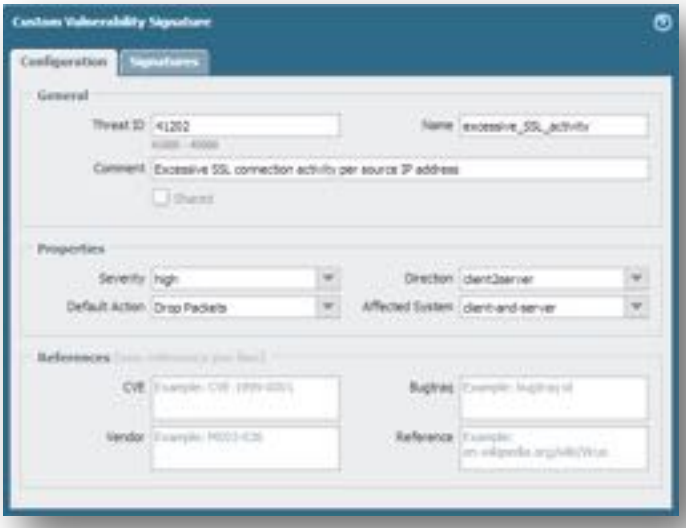
   Context 2 = http-req-uri-path
   Value = secure_portal (portal's path in URI string)



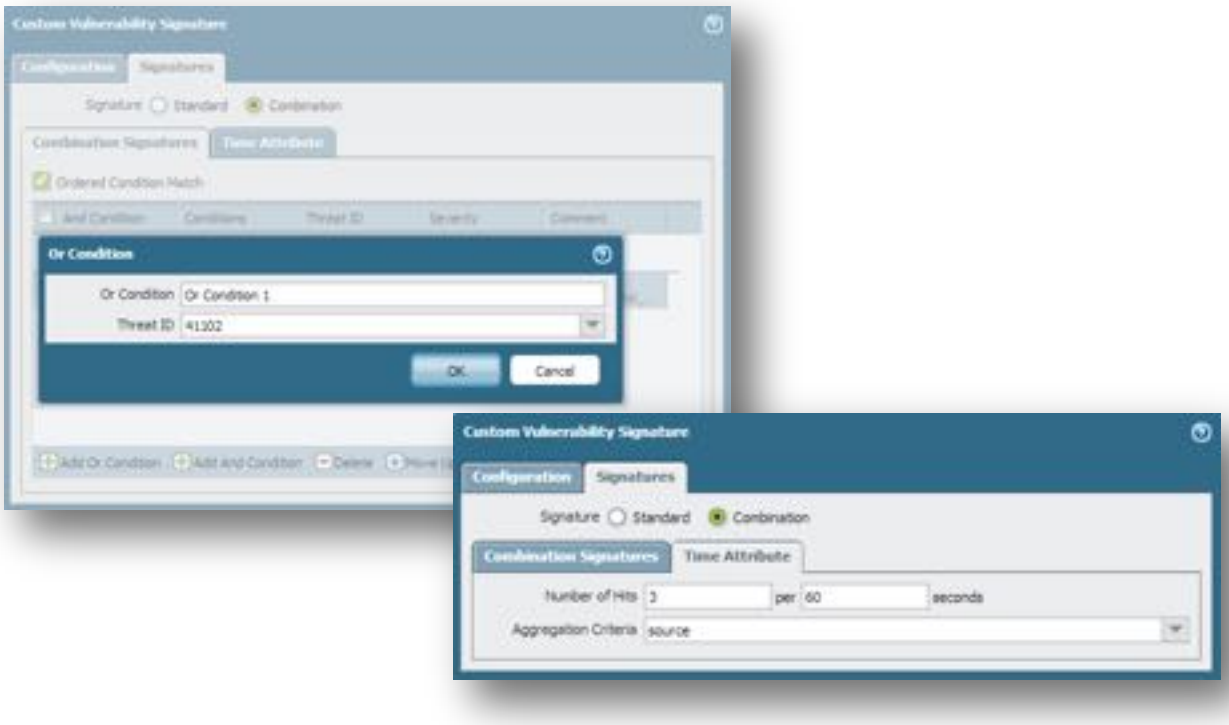
To summarize, this custom signature will look for SSL certificate requests with the string *secure_portal.mycompany.com* in the certificate's FQDN identifier. It will also match the *secure_portal* string from the URI path to ensure the certificate is coming from the secure portal. The *Transaction* scope setting allows the signature to match on individual transactions versus sessions to spot multiple SSL connection attempts within the same TCP session. When there is a match, a log event is created in the Threat Log with the *Alert* action with the *informational* severity.

2. Create a *combination signature* to combine the custom signature from the previous step and to set the time attributes for the protection threshold. In our example, well-behaved activity should require no more than two SSL connections in a 60 second interval. This example shows custom signature with Threat ID 41202 and the following characteristics:

   a. Configuration Tab defines the traffic & action properties:
      Severity = high
      Direction = client2server
      Default Action = Drop Packets
      Affected System = client-and-server

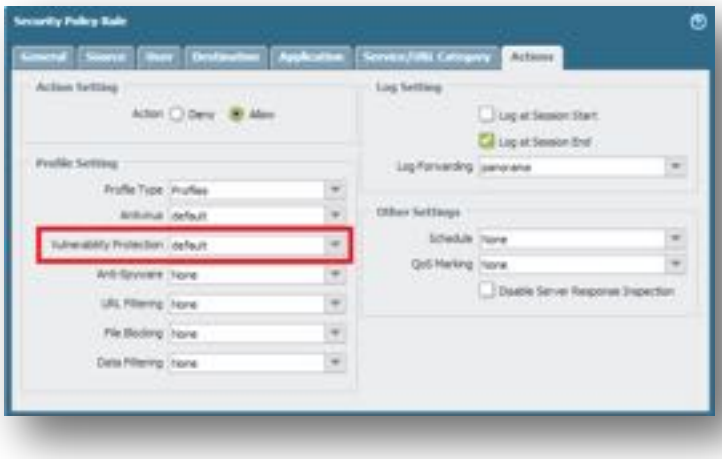


b. Signatures Tab defines the protection threshold:
   Signature = Combination
   Threat ID = specifies the custom vulnerability signature from step 1 (41102)
   Time Attribute = 3 (sets SSL connection threshold for 3 times per 60 second interval)
   Aggregation Criteria = source (from the client side)



To summarize, this combination signature will trigger when the same source IP address attempts more than three SSL connections within 60 seconds on the secure portal's web site using a certificate with the FQDN string specified in signature 41102. When the threshold is reached, the action will be to drop packets from the offending source IP address. An event with a *High* severity is logged in the Threat Log.

3. To activate the custom signatures, enable *Vulnerability Protection* on the security policy that controls access to the secure portal. All custom signatures are automatically enabled with the pre-defined (default and strict) vulnerability protection profiles.



4. If you need to bypass the custom vulnerability signatures for specific security policies, you can create a new vulnerability protection profile and exclude the custom signatures. See steps 3 and 5 from the first example for more details.

Excessive Search Attack

Having the ability to search for information is a valuable tool that nearly all web sites provide. Unfortunately, attackers can also use this great feature to attack the site by creating large volumes of search activity to exhaust the server's resources. Through the use of botnets, thousands of zombie PCs are instructed to perform random searches using various search strings. The database server(s) are flooded with thousands to millions of search requests over a short period of time, causing the servers to eventually run out of resources.

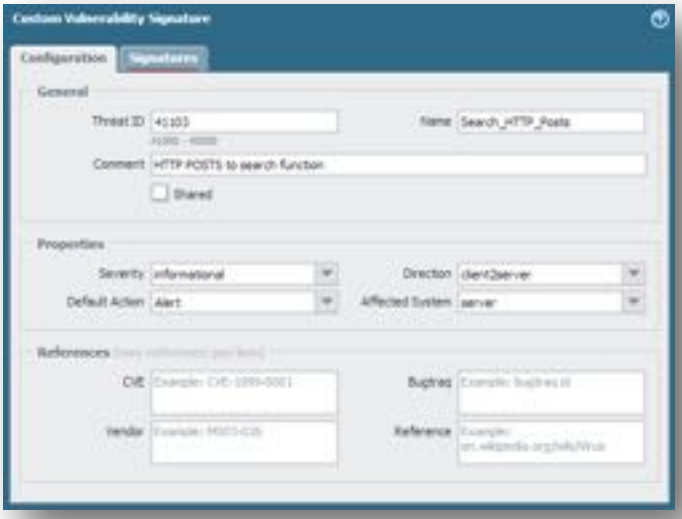
This example illustrates how a custom vulnerability signature can be created to monitor the number of HTTP POSTs to the search function on a specific web page. As it is perfectly legitimate for the clients to issue multiple searches in the course of using the website, the aggregation criteria for the search threshold will be performed on the server side. That is, the server's ability to process search requests from all clients will determine the protection threshold.

In this example, the company's IT team has already conducted the necessary capacity and performance tests to determine the maximum number of searches that can be supported by their database server(s). When the number of searches exceed the upper threshold set, the excess searches will be dropped.

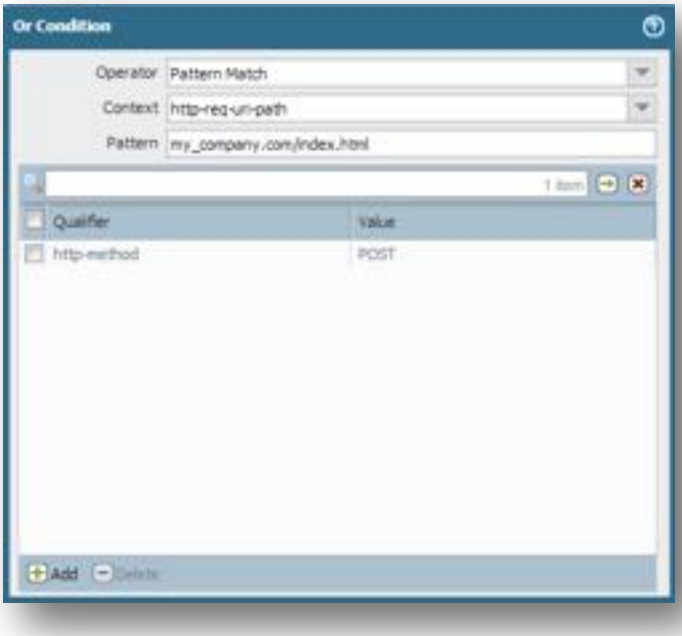
Deployment Example:

1. Create a custom vulnerability signature for the search application by matching on the HTTP POST operation (performing a search) and the string identifier for the URI path of the web site. This example shows custom signature with Threat ID 41103 and the following characteristics.

   a. Configuration Tab defines the traffic properties:
   Severity = informational
   Direction = client2server
   Default Action = Alert
   Affected System = server



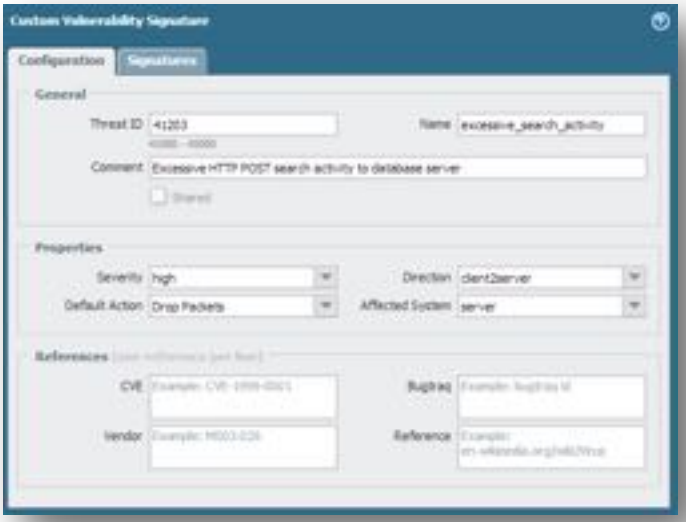
   b. Signatures Tab defines the match criteria:
   Scope = Transaction
   Operator = Pattern Match
   Context = http-req-uri-path
   Pattern = my_company.com/index.html (URI path)
   Qualifier = http-method for POST (search entry)



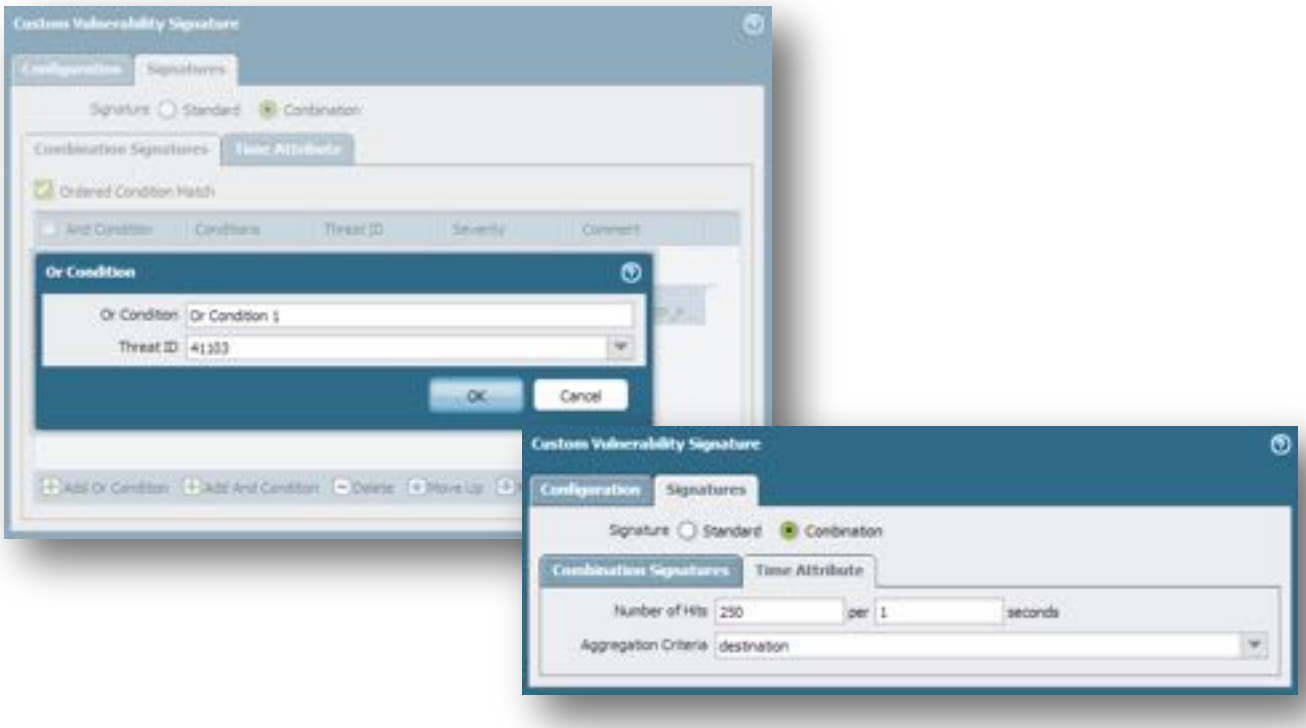
To summarize, this custom signature will look for *HTTP POST* activity to the company's home page which is identified by the *my_company.com/index.html* string in the URI path. The *Transaction* scope setting allows the signature to match on individual transactions versus sessions to spot multiple search transactions within the same TCP session. When there is a match, a log event is created in the Threat Log with the *Alert* action with the severity *informational*.

2. Create a *combination signature* to combine the custom signature from the previous step and to set the time attributes for the protection threshold. The company's database server performance and capacity testing helped to define the upper threshold - 250 searches per second before reaching 85% CPU utilization. The company would like to start limiting the searches once this threshold is reached. This example shows custom signature with Threat ID 41203 and the following characteristics:

   a. Configuration Tab defines the traffic & action properties:
      Severity = high
      Direction = client2server
      Default Action = Drop Packets
      Affected System = server



   b. Signatures Tab defines the protection threshold:
      Signature = Combination
      Threat ID = specifies the custom vulnerability signature from step 1 (41103)
      Time Attribute = 250 (sets HTTP POST 250 times per 1 second interval)
      Aggregation Criteria = destination (server side)

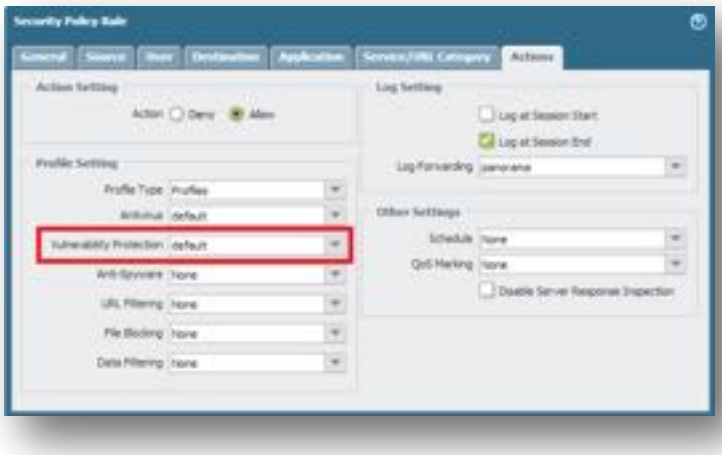


©2014, Palo Alto Networks, Inc.

To summarize, this combination signature will trigger when search volumes using HTTP POST exceed 250 per 1 second interval on the database server. When the threshold is reached, the action will be to drop packets to limit the search activity. An event with a *High* severity is logged in the Threat Log.

Note: Alternatively, the signature could have been configured to limit the number of search operations performed between EACH source IP address and the home page IP Address by using the *source-and-destination* aggregation criteria. Using this method would provide more granular limits for each source.

3. To activate the custom signatures, enable *Vulnerability Protection* on the security policy controlling access to the web site. All custom signatures are automatically enabled with the pre-defined (default and strict) vulnerability protection profiles.



4. If you need to bypass the custom vulnerability signatures for specific security policies, you can create a new vulnerability protection profile and exclude the custom signatures. See steps 3 and 5 from the first example for more details.

DNS Query Attack

PAN-OS's pre-defined vulnerability signatures provide a wide range of protection against DNS attacks (such as DNS ANY Brute Force, DNS Spoofing, DNS Amplification, DNS Zone Transfer, as well as many vendor specific DNS vulnerabilities). To view the latest DNS vulnerability signatures, search the Palo Alto Networks Threat Vault with the *DNS* search string as shown in the following illustration.



DNS is a common DoS attack vector because it is used widely in every network and the Internet. DNS cannot be blocked by firewalls without breaking applications and attackers rely on legitimate DNS functions to carry out their attacks. For example, an attack can be launched against the DNS infrastructure by requesting a name resolution for the same domain over and over again. When a botnet is used, thousands of zombie PCs can request millions of DNS resolution requests in a very short period of time and flood the DNS server or cause it to run out of resources.

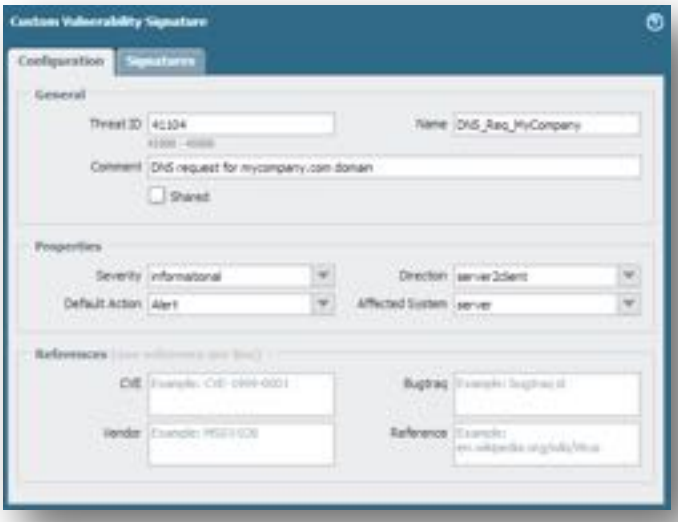
Modern operating systems and DNS servers will cache DNS responses by default, so there should be no reason why a host would repeatedly try to resolve the same domain name within a short period of time. In the case of many hosts behind a NAT'd network, they would most likely be using an internal DNS server that should be caching the DNS response from the authoritative server. So repeated resolutions for the same domain name originating from a single source IP address is abnormal behavior.

The following example illustrates how a custom signature is used to enforce the number of DNS requests from a single IP address for the same domain resolution.

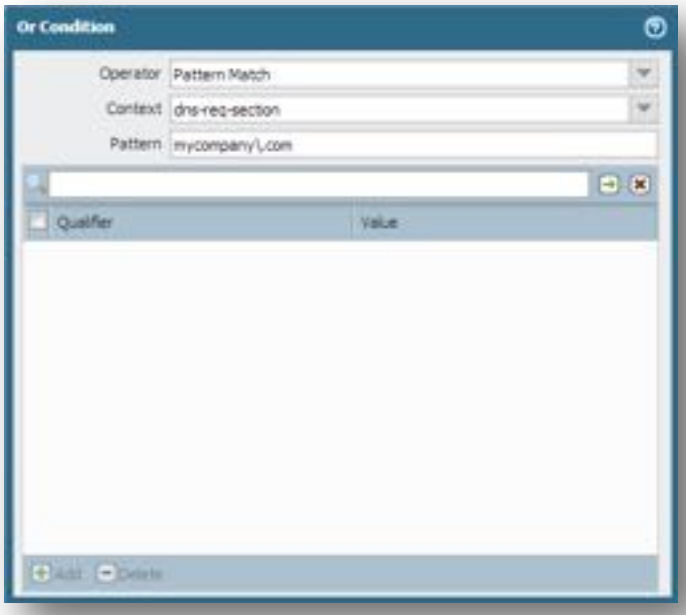
Deployment Example:

1. Create a custom vulnerability signature for the DNS request by matching on the domain name in the DNS request. This example shows custom signature with Threat ID 41104 and the following characteristics.

   a. Configuration Tab defines the traffic properties:
      Severity = informational
      Direction = client2server
      Default Action = Alert
      Affected System = server

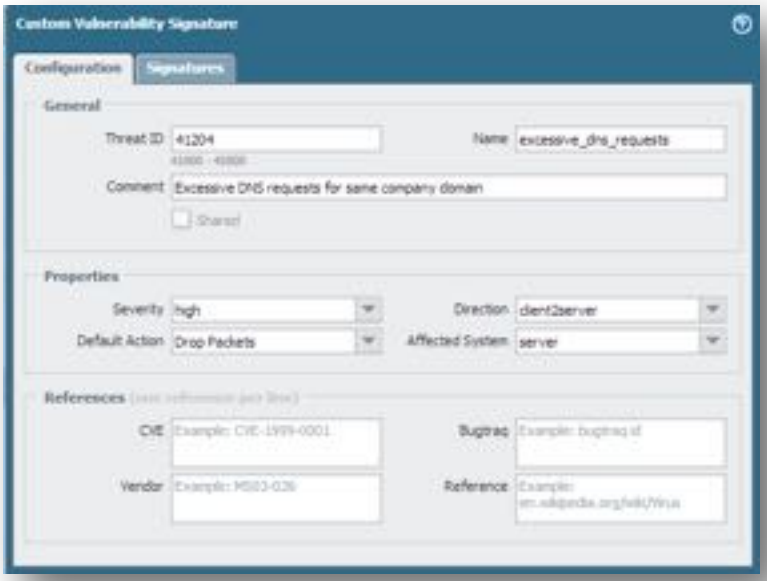


   b. Signatures Tab defines the match criteria:
      Scope = Session
      Operator = Pattern Match
      Context = dns-req-section
      Pattern = mycompany.com



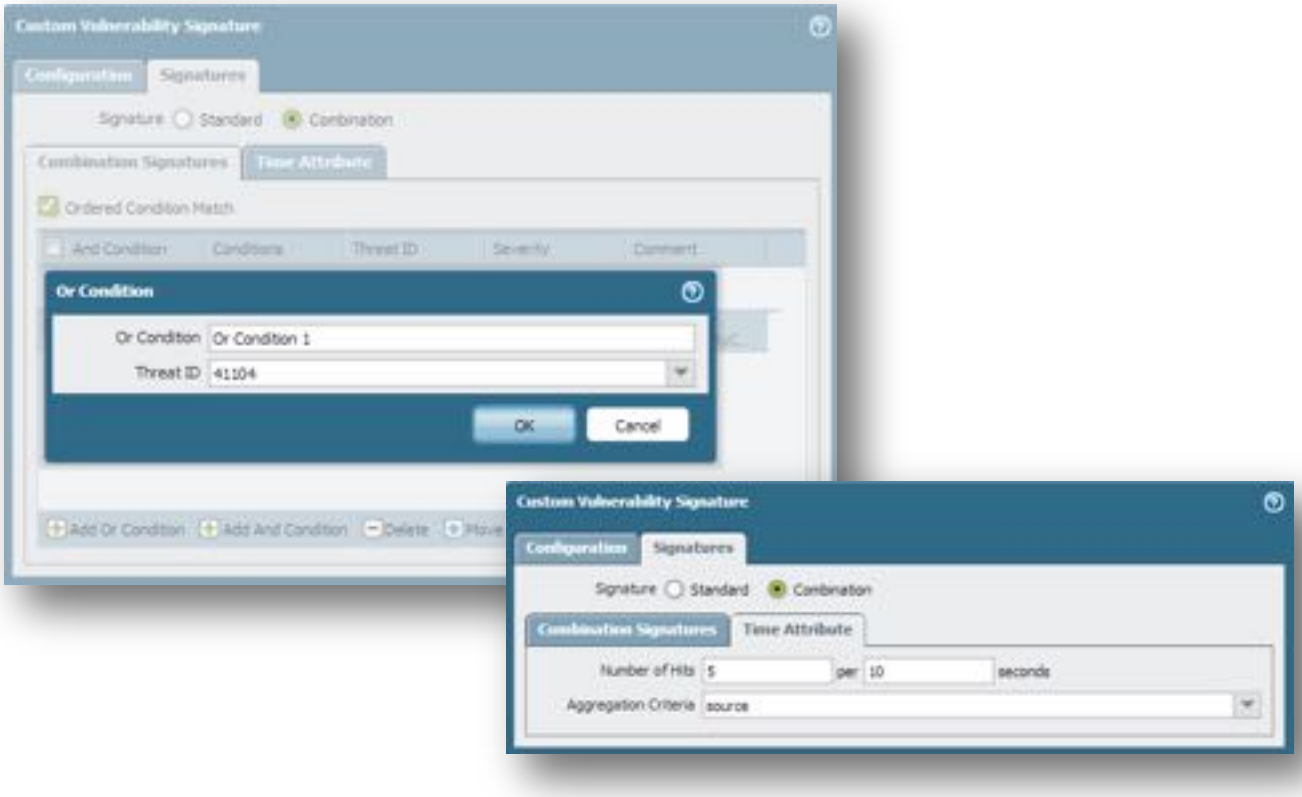
To summarize, this custom signature will look for the domain *mycompany.com* in a DNS request. The *Session* scope setting allows the signature to match on each DNS session from the client. When there is a match, a log event is created in the Threat Log with the *Alert* action at the *informational* severity.

2. Create a *combination signature* to combine the custom signature from the previous step and to set the time attributes for the protection threshold. This example shows custom signature with Threat ID 41204 and the following characteristics:



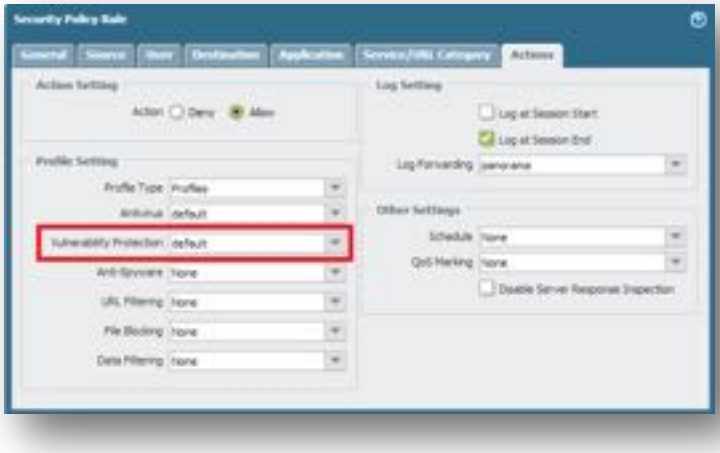
   a. Configuration Tab defines the traffic & action properties:
      Severity = high
      Direction = client2server
      Default Action = Drop Packets
      Affected System = server

   b. Signatures Tab defines the protection threshold:
      Signature = Combination
      Threat ID = specifies the custom vulnerability signature from step 1 (41104)
      Time Attribute = 5 (sets DNS request to 5 times per 10 second interval
      Aggregation Criteria = source (client side)



To summarize, this combination signature will trigger when the number of DNS requests for *mycompany.com* exceeds 5 within a 10 second interval from a single source address. When the threshold is reached, the action will be to drop packets to limit the DNS activity. An event with a *High* severity is logged in the Threat Log.

3. To activate the custom signatures, enable *Vulnerability Protection* on the security policy controlling access to the DNS server. All custom signatures are automatically enabled with the pre-defined (default and strict) vulnerability protection profiles.

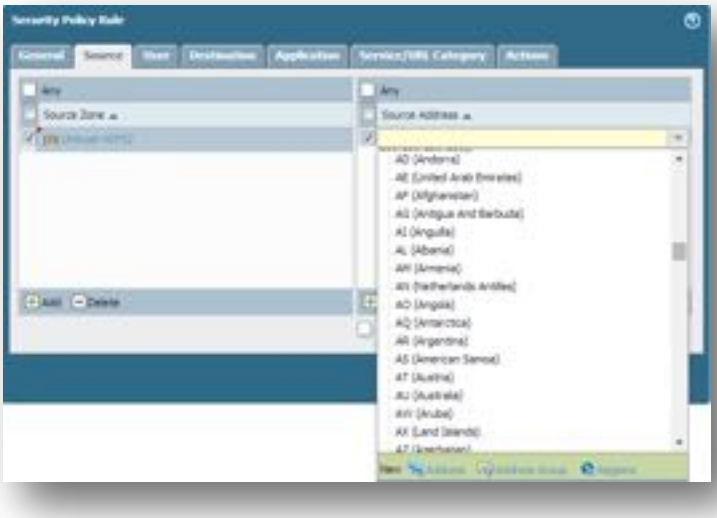


4. If you need to bypass the custom vulnerability signatures for specific security policies, you can create a new vulnerability protection profile and exclude the custom signatures. See steps 3 and 5 from the first example for more details.

Mitigation with Geo Locations

Having the ability to limit or restrict traffic based on its origin can be an effective way to reduce the attack surface of an application. For example, if a web service is only designed for its members who are living in Germany, granting access to German IP addresses and limiting all others can be an effective way to reduce application attacks from regions outside of Germany. However, blocking all non-German IP addresses may create false positives as the members may be travelling abroad and require access.

With PAN-OS security policies and custom vulnerability signatures, polices can be created to allow full access to specific geo locations and limited or controlled access to other regions.



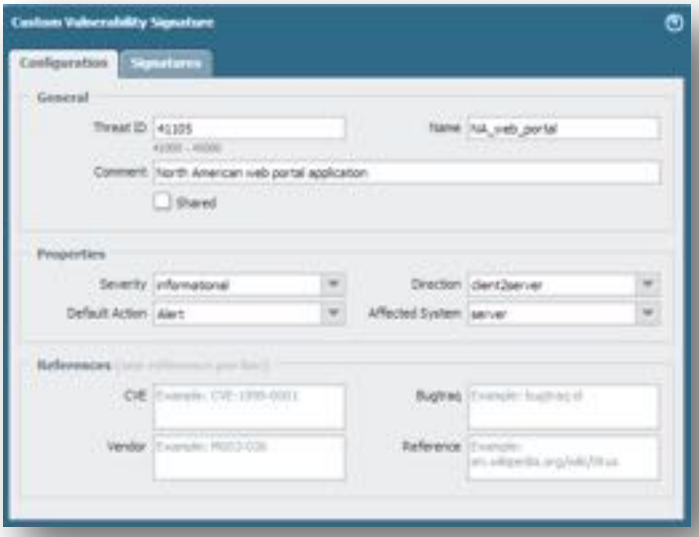
Specific geo locations can be specified in the security policy's *Source Address* filter.

This example illustrates how custom vulnerability signatures can be used with geo location specific security policies to reduce the risk of attacks. A custom vulnerability signature will be created for a web site's home page that is designed specifically for the North American audience and will be combined with a security policy that restricts access based on geo location. Customers who are physically located inside North America will not be restricted, but access outside North America will be controlled by throttling the number of connections to the web site.

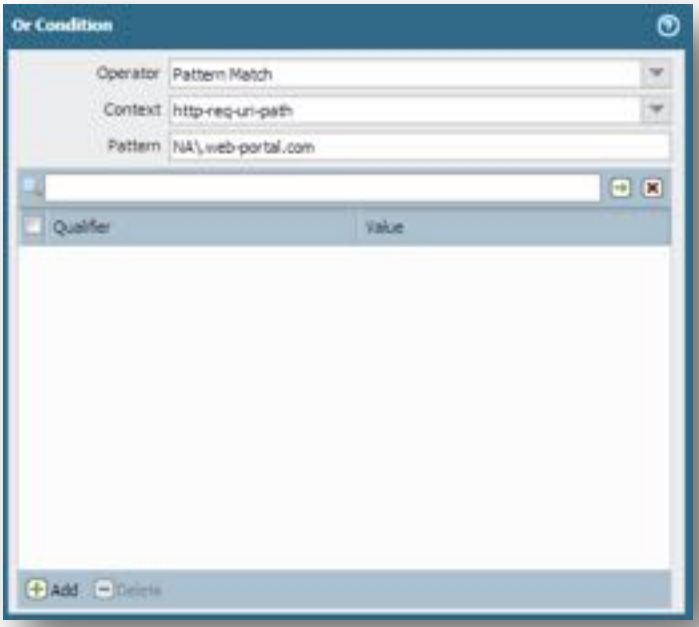
Deployment Example:

1. Create a custom vulnerability signature for the web application by matching on the URI string of its home page. This example shows custom signature with Threat ID 41105 and the following characteristics.

   a. Configuration Tab defines the traffic properties:
      Severity = informational
      Direction = client2server
      Default Action = Alert
      Affected System = server



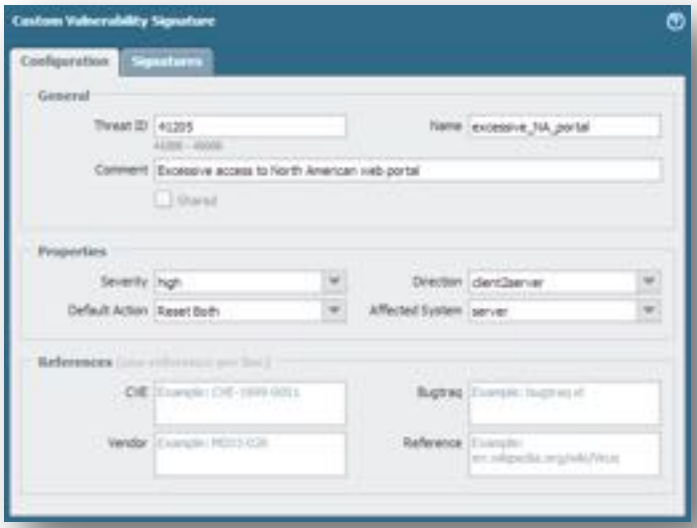
   b. Signatures Tab defines the match criteria:
      Scope = Session
      Operator = Pattern Match
      Context = http-req-uri-path
      Pattern = NA.web-portal.com



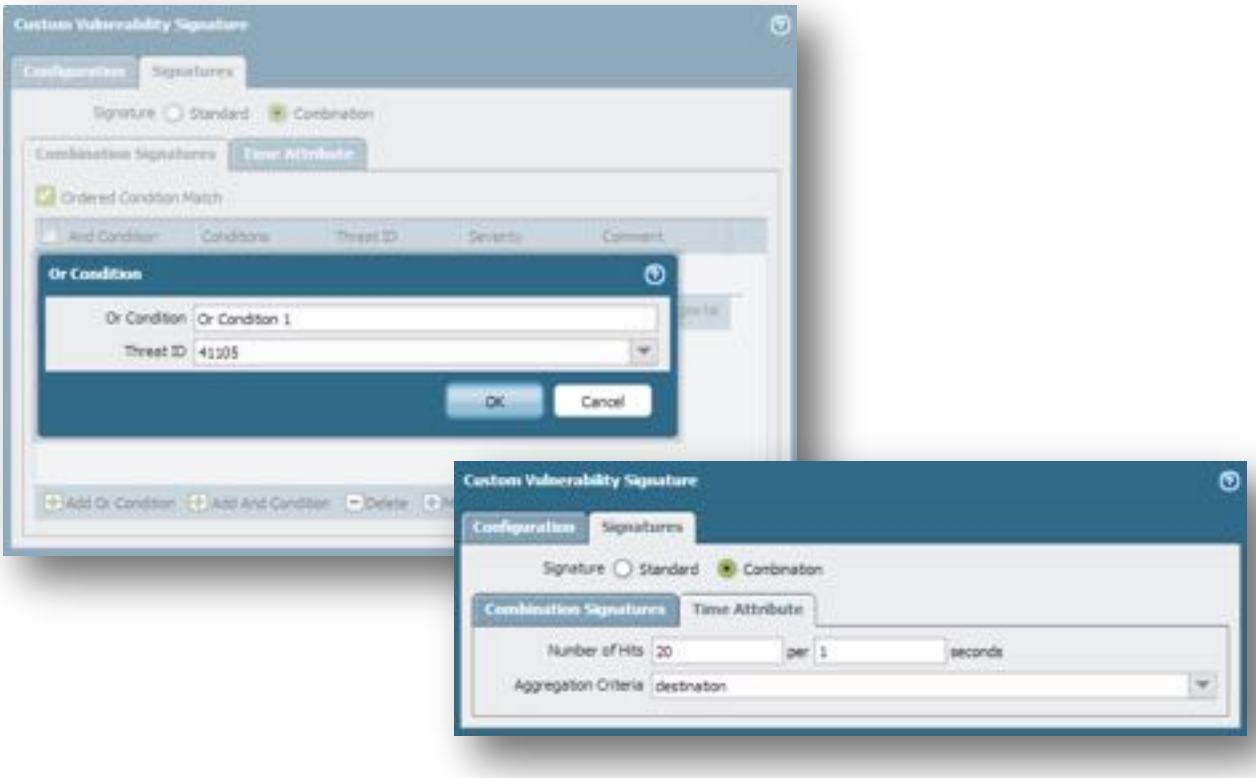
   To summarize, this custom signature will look for the string *NA.web-portal.com* in the URI path. The *Session* scope setting allows the signature to match on each session from the client. When there is a match, a log event is created in the Threat Log with the *Alert* action at the *informational* severity.

2. Create a *combination signature* to combine the custom signature from the previous step and to set the time attributes for the protection threshold. The company will only permit non-North American access at a rate of 20 sessions per second for the site's server. This allows the web application to be accessed globally, but at a much lower rate that will never be able to DoS the site and false positives for its North American customers travelling abroad are eliminated. This example shows custom signature with Threat ID 41205 and the following characteristics:

a. Configuration Tab defines the traffic & action properties:
   Severity = high
   Direction = client2server
   Default Action = reset-both
   Affected System = server



b. Signatures Tab defines the protection threshold:
   Signature = Combination
   Threat ID = specifies the custom vulnerability signature from step 1 (41105)
   Time Attribute = 20 (sets access to 20 times per 1 second interval)
   Aggregation Criteria = destination (server side)



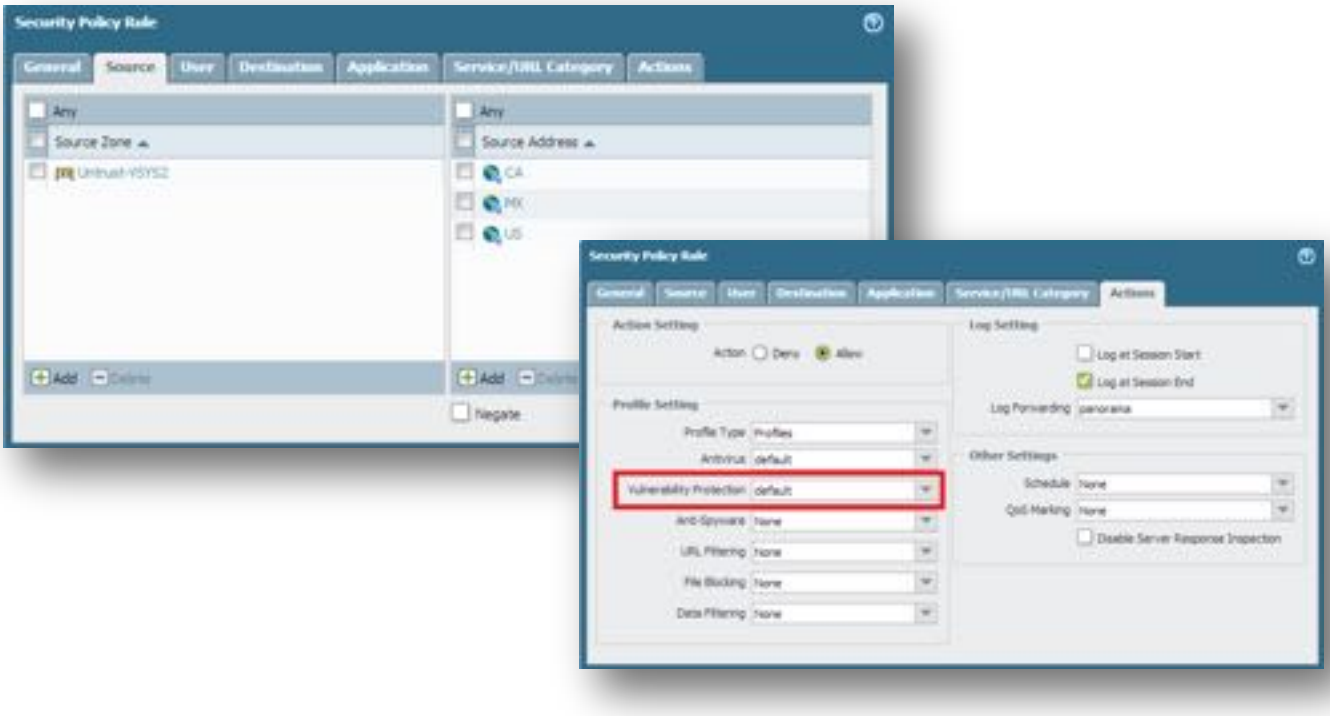
To summarize, this combination signature will trigger when the number of sessions containing the *NA.web-portal.com* string in the URI path is over 20 sessions per second on the server. When the threshold is reached, the action will be to reset both the client and the server to end the session gracefully. An event with a *High* severity is logged in the Threat Log.

©2014, Palo Alto Networks, Inc.

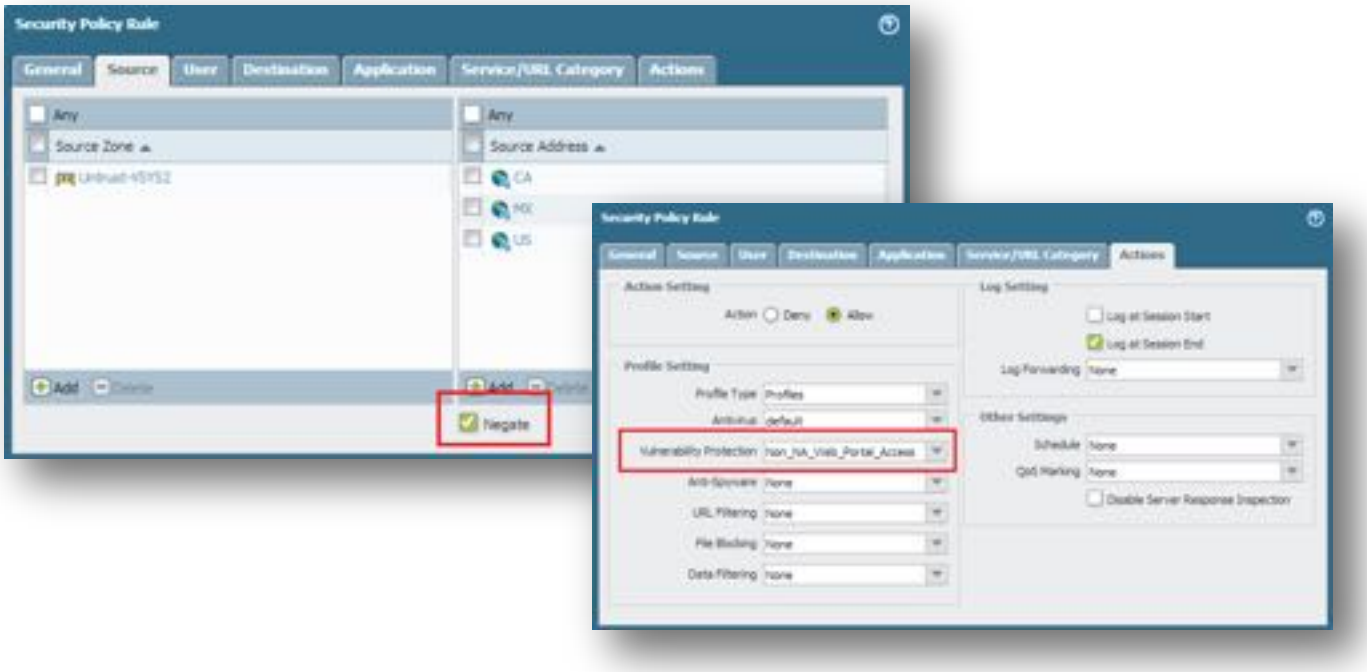
3. To activate the custom signatures, enable *Vulnerability Protection* on the security policy controlling access to the North American web site. Two security policies will be required to control access.
    a. The first security policy will match only North American geo locations (Canada, United States, and Mexico) and will not have a vulnerability profile that limits access capacity to the web site.
    b. The second security profile will match all non-North American geo locations (with Negate option) and will contain the vulnerability profile that limits access capacity as defined in the custom vulnerability signatures.



This illustration shows the security policy that will apply the default vulnerability protection policy (without the access limitations) to the North American geo locations.



This next illustration shows the security policy that applies the vulnerability protection profile (with the access limits) to all other geo locations except for the North American countries. Note that the *Negate* checkbox is used to quickly select all other geo locations that do NOT match the ones selected.



Summary

Custom vulnerability signatures can be used to provide granular application visibility. When used with combination signatures, they can be a powerful tool to spot abnormal application use and provide enterprises with the means to proactively mitigate application-based DoS attacks. The use cases shown in this paper are just a few examples of modern application DoS attacks and the examples are designed to show the flexibility of custom vulnerability signatures and how they can be used to provide better visibility and control against these types of attacks.

Revision History

Date	Revision	Comment
June 12, 2014	B	

©2014, Palo Alto Networks, Inc.

Palo Alto Networks® Compatibility Matrix

techDOCS paloalto NETWORKS

Contact Information

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

About the Documentation

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

Copyright

Palo Alto Networks, Inc.
www.paloaltonetworks.com

© 2016-2018 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

Last Revised

November 1, 2018

Table of Contents

Supported OS Releases by Model 5
Palo Alto Networks Next-Gen Firewalls 7
Palo Alto Networks Appliances 9
HA Port and Processor Support 10

VM-Series Firewalls 13
VM-Series Firewall Hypervisor Support 15
Partner Interoperability for VM-Series Firewalls 21
AWS and AWS Gov Cloud Regions 22
Azure Regions 23
Google Cloud Regions 24
AWS CFT Amazon Machine Images (AMI) List 25
Images for PAN-OS 8.1 25
Images for PAN-OS 8.0 26
Images for PAN-OS 7.1.11 29

Panorama 31
Plugins 33
VMware NSX 33
Cloud Services 33
Interconnect 34
Public Cloud-AWS, Azure, GCP 34
Panorama Hypervisor Support 36

MFA Vendor Support 39
MFA Vendor Support 41

Supported Cipher Suites 43
Cipher Suites Supported in PAN-OS 8.1 45
PAN-OS 8.1 GlobalProtect Cipher Suites 45
PAN-OS 8.1 IPSec Cipher Suites 46
PAN-OS 8.1 IKE and Web Certificate Cipher Suites 47
PAN-OS 8.1 Decryption Cipher Suites 48
PAN-OS 8.1 Administrative Session Cipher Suites 50
PAN-OS 8.1 HA1 SSH Cipher Suites 51
PAN-OS 8.1 PAN-OS-to-Panorama Connection Cipher Suites 52
PAN-OS 8.1 Cipher Suites Supported in FIPS-CC Mode 52
Cipher Suites Supported in PAN-OS 8.0 55
PAN-OS 8.0 GlobalProtect Cipher Suites 55
PAN-OS 8.0 IPSec Cipher Suites 56
PAN-OS 8.0 IKE and Web Certificate Cipher Suites 57
PAN-OS 8.0 Decryption Cipher Suites 58
PAN-OS 8.0 Administrative Session Cipher Suites 60
PAN-OS 8.0 PAN-OS-to-Panorama Connection Cipher Suites 61
PAN-OS 8.0 Cipher Suites Supported in FIPS-CC Mode 62
Cipher Suites Supported in PAN-OS 7.1 65

PAN-OS 7.1 GlobalProtect Cipher Suites....65
PAN-OS 7.1 IPSec Cipher Suites....66
PAN-OS 7.1 IKE and Web Certificate Cipher Suites....67
PAN-OS 7.1 Decryption Cipher Suites....68
PAN-OS 7.1 Administrative Session Cipher Suites....69
PAN-OS 7.1 PAN-OS-to-Panorama Connection Cipher Suites....70

GlobalProtect....73
Where Can I Install the GlobalProtect App?....75
What X-Auth IPSec Clients are Supported?....78
What Features Does GlobalProtect Support?....79
What GlobalProtect Features Do Third-Party Mobile Device Management Systems Support?....83

User-ID Agent....85
Where Can I Install the User-ID Agent?....87
Which Servers Can the User-ID Agent Monitor?....88

Terminal Services (TS) Agent....89
Where Can I Install the Terminal Services (TS) Agent?....91
How Many TS Agents Does My Firewall Model Support?....92

Traps....93
Where Can I Install the Endpoint Security Manager (ESM)?....95
Where Can I Install the Traps Agent?....96
Endpoint Operating Systems Supported with Traps....96
Virtual Desktop Operating Systems Supported with Traps....101
Virtual Applications Supported with Traps....101

IPv6 Support by Feature....103
IPv6 Support by Feature....105

Supported OS Releases by Model

Use the tables throughout this Palo Alto Networks® Compatibility Matrix to determine support for Palo Alto Networks next-generation firewalls, appliances, and agents. Additionally, refer to the product comparison tool for detailed information about Palo Alto Networks firewalls by model, including specifications for throughput, maximum number of sessions, rules, objects, tunnels, and zones.

For supported operating systems on firewalls and appliances and for high-availability (HA) port and processor support on firewalls, review the following topics:

- Palo Alto Networks Next-Gen Firewalls
- Palo Alto Networks Appliances
- HA Port and Processor Support

© 2018 Palo Alto Networks, Inc.

Palo Alto Networks Next-Gen Firewalls

The following table shows the PAN-OS® releases supported for each firewall model.

Palo Alto Networks Firewall Model	PAN-OS 6.1	PAN-OS 7.1	PAN-OS 8.0	PAN-OS 8.1
Hardware Firewalls				
PA-200 Firewall (EoS*)	✓	✓	✓	✓
PA-220 Firewall	–	–	✓	✓
PA-220R Firewall	–	–	–	✓
PA-500 Firewall (EoS*)	✓	✓	✓	✓
PA-800 Series Firewall	–	–	✓	✓
PA-2000 Series Firewall (EoS*)	✓	✓	–	–
PA-3000 Series Firewall	✓	✓	✓	✓
PA-3200 Series Firewall	–	–	–	✓
PA-4000 Series Firewall (EoS*)	✓	✓	–	–
PA-5000 Series Firewall	✓	✓	✓	✓
PA-5220, PA-5250, and PA-5260 Firewalls	–	–	✓	✓
PA-5280 Firewall	–	–	–	✓
PA-7050 Firewall	✓	✓	✓	✓
PA-7080 Firewall	–	✓	✓	✓
VM-Series Firewalls				
VM-50 Firewall	–	–	✓	✓
VM-100 Firewall	✓	✓	✓	✓
VM-200 Firewall	✓	✓	✓	✓
VM-300 Firewall	✓	✓	✓	✓

© 2018 Palo Alto Networks, Inc.

Palo Alto Networks Firewall Model	PAN-OS 6.1	PAN-OS 7.1	PAN-OS 8.0	PAN-OS 8.1
VM-500 Firewall	–	–	✓	✓
VM-700 Firewall	–	–	✓	✓
VM-1000-HV Firewall	✓	✓	✓	✓

* For more specific information about firewalls and appliances that have reached end-of-sale (EoS) status, refer to Hardware End-of-Life Dates.

© 2018 Palo Alto Networks, Inc.

Palo Alto Networks Appliances

The following table shows OS release support for each Palo Alto Networks appliance.

Palo Alto Networks Appliance	Release 6.1	Release 6.2	Release 7.1	Release 8.0	Release 8.1
GP-100 Appliance (EoS*)	✓	✓	–	–	–
Panorama Virtual Appliance	✓	–	✓	✓ 8.0.2 & later	✓
M-100 Appliance (EoS*)	✓	–	✓	✓ 8.0.2 & later	✓
M-200 Appliance	–	–	–	–	✓
M-500 Appliance	–	–	✓	✓ 8.0.2 & later	✓
M-600 Appliance	–	–	–	–	✓
WF-500 Appliance	✓	–	✓	✓ 8.0.1 & later	✓

* For more specific information about firewalls and appliances that have reached end-of-sale (EoS) status, refer to Hardware End-of-Life Dates.

© 2018 Palo Alto Networks, Inc.

HA Port and Processor Support

The following table identifies which Palo Alto Networks® next-gen firewalls can support the HA ports and processor functionality you require in your network.

Palo Alto Networks Firewall Model	Separate Mgmt Plane Processor	Network Processor	Signature Match	First Packet Processor	HA1 Port	HA2 Port	HSCI Port
PA-200 (EoS*)	–	–	–	–	–	–	–
PA-220	–	–	–	–	–	–	–
PA-220R	–	–	–	–	–	–	–
PA-500 (EoS*)	✓	–	–	–	–	–	–
PA-820	–	–	–	–	✓	✓	–
PA-850	–	–	–	–	✓	✓	–
PA-3020	✓	–	✓	–	✓	✓	–
PA-3050	✓	✓	✓	–	✓	✓	–
PA-3060	✓	✓	✓	–	✓	✓	–
PA-3220	✓	✓	–	–	✓ (x2)	–	✓
PA-3250	✓	✓	✓	–	✓ (x2)	–	✓
PA-3260	✓	✓	✓	–	✓ (x2)	–	✓
PA-5020	✓	✓	✓	–	✓	✓	–
PA-5050	✓	✓	✓	–	✓	✓	–
PA-5060	✓	✓	✓	–	✓	✓	–
PA-5220	✓	✓	✓	✓	✓ (x2)	–	✓
PA-5250	✓	✓	✓	✓	✓ (x2)	–	✓
PA-5260	✓	✓	✓	✓	✓ (x2)	–	✓

© 2018 Palo Alto Networks, Inc.

Palo Alto Networks Firewall Model	Separate Mgmt Plane Processor	Network Processor	Signature Match	First Packet Processor	HA1 Port	HA2 Port	HSCI Port
PA-5280	✓	✓	✓	✓	✓ (x2)	–	✓
PA-7050	✓	✓	✓	✓	✓ (x2)	–	✓ (x2)
PA-7080	✓	✓	✓	✓	✓ (x2)	–	✓ (x2)

* For more specific information about firewalls and appliances that have reached end-of-sale (EoS) status, refer to Hardware End-of-Life Dates.

© 2018 Palo Alto Networks, Inc.

VM-Series Firewalls

The hypervisors and the public cloud regions in which you can deploy the VM-Series firewalls:

- VM-Series Firewall Hypervisor Support
- Partner Interoperability for VM-Series Firewalls
- AWS and AWS Gov Cloud Regions
- Azure Regions
- Google Cloud Regions
- AWS CFT Amazon Machine Images (AMI) List

© 2018 Palo Alto Networks, Inc.

VM-Series Firewall Hypervisor Support

The following table shows hypervisor version support on the VM-Series firewall.

Deployment	PAN-OS Version Support	Hypervisor Version Support	I/O Enhancement Support	Base Image Required from the Palo Alto Networks Support Portal
Private Cloud Deployments				
VM-Series for VMware vSphere Hypervisor (ESXi) (without VMware NSX)	PAN-OS 6.1	5.0, 5.1, 5.5	—	**PAN-OS for VM-Series Base Images** For example, the download-able image name is: **PA-VM-6.1.0.zip**
	PAN-OS 7.1 PAN-OS 8.0 PAN-OS 8.1	5.1, 5.5, 6.0	—	**PA-VM-ESX-7.1.0.ova** **PA-VM-ESX-8.0.0.ova** **PA-VM-ESX-8.1.0.ova**
	PAN-OS 7.1 PAN-OS 8.0 PAN-OS 8.1	6.5	PAN-OS 8.0 and later: • SR-IOV • DPDK	**PA-VM-ESX-7.1.0-u1.ova** **PA-VM-ESX-8.0.0.ova** **PA-VM-ESX-8.1.0.ova**
VM-Series for VMware NSX vSphere with VMware NSX and Panorama The combinations listed in this document have been approved by Palo Alto Networks. For versions of PAN-OS certified by VMware, see the VMware Compatibility Guide.	PAN-OS 6.1	vSphere: 5.5 NSX Manager: 6.1	—	**PAN-OS for VM-Series NSX Base Images** For example, the download-able image name is: **PA-VM-NSX-6.1.0.zip**
	PAN-OS 7.1	vSphere: 5.5 NSX Manager: 6.0, 6.1, 6.2	—	**PA-VM-NSX-7.1.0.zip**
		vSphere: 6.0, 6.5a NSX Manager: 6.1, 6.2, 6.3	—	

© 2018 Palo Alto Networks, Inc.

Deployment	PAN-OS Version Support	Hypervisor Version Support	I/O Enhancement Support	Base Image Required from the Palo Alto Networks Support Portal
Panorama 8.0 and higher requires the VMware NSX plugin. For more information, see VMware NSX. Available on all VM-Series models except the VM-50 and the VM-700.	PAN-OS 8.0 with NSX Plugin 1.0.0	vSphere: 5.5, 6.0 NSX Manager: 6.1, 6.2	LRO	**PA-VM-NSX-8.0.0.zip** with **vmware_nsx-1.0.0**
	PAN-OS 8.0 with NSX Plugin 1.0.1 or later	vSphere: 5.5, 6.0, NSX Manager: 6.1, 6.2, 6.3	LRO	**PA-VM-NSX-8.0.0.zip** with **vmware_nsx-1.0.1**
		vSphere: 6.5a NSX Manager: 6.3		
	PAN-OS 8.0.5 with NSX Plugin 2.0.2 or later	vSphere: 6.0, 6.5 NSX Manager: 6.3.x and 6.4.0	LRO	**PA-VM-NSX-8.0.5.zip** with **vmware_nsx-2.0.2**
		vSphere: 6.0, 6.5, 6.7 NSX Manager: 6.4.1 or later		
	PAN-OS 8.1 with NSX Plugin 2.0.2 or later	vSphere: 6.0, 6.5 NSX Manager: 6.3.x, 6.4.0	LRO	**PA-VM-NSX-8.1.0.zip** with **vmware_nsx-2.0.2**
		vSphere: 6.0, 6.5,6.7 NSX Manager: 6.4.1 or later		
VM-Series for Citrix SDX	PAN-OS 6.1 PAN-OS 7.1 PAN-OS 8.0	SDX 10.1+ XenServer 6.0.2 or later	—	**PAN-OS for VM-Series SDX Base Images** For example, the download-able image name is: **PA-VM-SDX-6.1.0.zip** **PA-VM-SDX-7.1.2.zip** **PA-VM-SDX-8.0.0.zip**
VM-Series for KVM	PAN-OS 6.1	KVM on the following Linux distributions:	SR-IOV	**PAN-OS for VM-Series KVM Base Images**

© 2018 Palo Alto Networks, Inc.

Deployment	PAN-OS Version Support	Hypervisor Version Support	I/O Enhancement Support	Base Image Required from the Palo Alto Networks Support Portal
		• Ubuntu: 12.04 LTS, 14.04 LTS • CentOS/ Red Hat Enterprise Linux: 6.5, 7.0, 7.1		For example, the download-able image name is: **PA-VM-6.1.0.qcow2**
	PAN-OS 7.1	• Ubuntu: 12.04 LTS, 14.04 LTS, 16.04 LTS • CentOS/ Red Hat Enterprise Linux: 6.5, 7.0, 7.1, 7.2	SR-IOV	**PA-VM-7.1.0.qcow2**
	PAN-OS 8.0	• Ubuntu: 14.04 LTS, 16.04 LTS • CentOS/ Red Hat Enterprise Linux: 7.2	• DPDK • SR-IOV	**PA-VM-8.0.0.qcow2**
	PAN-OS 8.1	• Ubuntu: 14.04 LTS, 16.04 LTS • CentOS/ Red Hat Enterprise Linux: 7.2	• DPDK • SR-IOV	**PA-VM-8.1.0.qcow2**
	PAN-OS 8.1.3	CentOS: 7.5	Packet MMAP mode support only; so you must disable DPDK with op-cmd-dpdk-pkt-io=off in the init-cfg.txt file used for bootstrapping or use the CLI command set system setting dpdk-pkt-io off)	**PA-VM-8.1.3.qcow2**

© 2018 Palo Alto Networks, Inc.

Deployment	PAN-OS Version Support	Hypervisor Version Support	I/O Enhancement Support	Base Image Required from the Palo Alto Networks Support Portal
VM-Series for Nutanix	PAN-OS 8.1.0 Layer 2, Layer 3, Virtual wire deployments	AHV Nutanix AOS Version - 5.1.5 Nutanix AHV Release 20160925.121	—	**PA-VM-8.1.0.qcow2**
	PAN-OS 8.1.2 Layer 3 deployments only	AHV Nutanix AOS Version - 5.5.4 or later Nutanix AHV Release 20170830.156S	Packet MMAP support for Layer 3 deployment. You must disable DPDK with **`op-cmd-dpdk-pkt-io=off`** in the init-cfg.txt file used for bootstrapping or use the CLI command **`set system setting dpdk-pkt-io off`**.	**PA-VM-8.1.0.qcow2** and upgrade to PAN-OS 8.1.2.
	PAN-OS 8.1.3 Layer 3 deployments, and Virtual wire deployments with Service Chaining.	AHV Nutanix AOS Version - 5.9 Nutanix AHV Release 20170830.171	Packet MMAP support only. You must disable DPDK with **`op-cmd-dpdk-pkt-io=off`** in the init-cfg.txt file used for bootstrapping or use the CLI command **`set system setting dpdk-pkt-io off`**.	**PA-VM-8.1.3.qcow2**
VM-Series for Hyper-V	PAN-OS 7.1 PAN-OS 8.0	• Windows Server 2012 R2 with Hyper-V role or Hyper-V 2012 R2	—	**PAN-OS for VM-Series Hyper-V Base Images** For example, the download-able image name is:

© 2018 Palo Alto Networks, Inc.

Deployment	PAN-OS Version Support	Hypervisor Version Support	I/O Enhancement Support	Base Image Required from the Palo Alto Networks Support Portal
	PAN-OS 8.1	• Windows Server 2016 with Hyper-V role or Hyper-V 2016		**PA-VM-HPV-7.1.0.vhdx** **PA-VM-HPV-8.0.0.vhdx** **PA-VM-HPV-8.1.0.vhdx**
VM-Series for OpenStack	PAN-OS 7.1.12 PAN-OS 8.0 PAN-OS 8.1	KVM: Ubuntu 14.04 OpenStack: Mirantis 8.0 with Contrail 3.2	—	**PAN-OS for VM-Series KVM Base Images** PA-VM-7.1.12.qcow2 PA-VM-8.0.0.qcow2 PA-VM-8.1.0.qcow2
Hardware and VM-Series Firewalls in Cisco ACI	PAN-OS 7.1 PAN-OS 8.0	ACI: 2.3, 3.0	N/A	Device Package 1.3
Public Cloud Deployments				
VM-Series on AWS List of supported AWS and AWS Gov Cloud Regions.	PAN-OS 6.1 PAN-OS 7.1 PAN-OS 8.0 PAN-OS 8.1	N/A	PAN-OS 8.0 and later:DPDK is supported, but disabled by default	N/A
VM-Series on vCloud Air	PAN-OS 7.1 PAN-OS 8.0 PAN-OS 8.1	N/A	—	**PAN-OS for VM-Series Base Images** For example, the download-able image name is: PA-VM-ESX-7.1.0.ova PA-VM-ESX-8.0.0.ova PA-VM-ESX-8.1.0.ova
VM-Series on Azure List of supported Azure Regions.	PAN-OS 7.1 PAN-OS 8.0 PAN-OS 8.1	N/A	DPDK is not supported	N/A

© 2018 Palo Alto Networks, Inc.

Deployment	PAN-OS Version Support	Hypervisor Version Support	I/O Enhancement Support	Base Image Required from the Palo Alto Networks Support Portal
VM-Series on Google Cloud	PAN-OS 8.1	N/A	DPDK is supported and enabled by default	N/A

© 2018 Palo Alto Networks, Inc.

Partner Interoperability for VM-Series Firewalls

The following table shows the partner products with which the VM-Series firewall interoperates.

Partner Product	Partner Hardwa	Hypervisor	Partner Software Version	PAN-OS Version	Deployment Modes Supported	Documentation
Juniper NFX Network Services Platform	NFX 250	KVM	Junos 15.1X53-D470.5	PAN-OS 8.0.5 PAN-OS 8.1.0	Layer 2, Layer3, Virtual wire deployments VirtIO only	Juniper NFX documentation
NSX SD-WAN by VeloCloud	Edge 520v Edge 840	KVM	VCE 3.2.0	PAN-OS 8.0.5	Virtual wire deployments VirtIO only	NSX SD-WAN by VeloCloud documentation
Nutanix	—	AHV	See VM-Series on Nutanix.			Nutanix documentation
Cisco Enterprise Network Compute System(ENCS)	Cisco 5400 Series	KVM	NFVIS version 3.8.1	PAN-OS 8.1.3	Layer 2, Layer3, Virtual wire deployments VirtIO only	VM-Series on Cisco ENCS

© 2018 Palo Alto Networks, Inc.

AWS and AWS Gov Cloud Regions

The AWS regions in which you can deploy the VM-Series firewall from the AWS Marketplace.

AWS Regions	Region ID
US East (N. Virginia)	us-east-1
US East (Ohio)	us-east-2
US West (N.California)	us-west-1
US West (Oregon)	us-west-2
Asia Pacific (Singapore)	ap-southeast-1
Asia Pacific (Sydney)	ap-southeast-2
Asia Pacific (Tokyo)	ap-northeast-1
Asia Pacific (Seoul)	ap-northeast-2
Asia Pacific (Mumbai)	ap-south-1
Canada Central	ca-canada-1
EU (Frankfurt)	eu-central-1
EU (Ireland)	eu-west-1
EU (London)	eu-west-2
South America (Sao Paulo)	sa-east-1
AWS Gov Cloud (US)	us-gov-west
Asia Pacific (Beijing) (shared AMI, no marketplace)	AWS China

© 2018 Palo Alto Networks, Inc.

Azure Regions

The Azure locations in which you can deploy the VM-Series firewall from the Azure public and the Azure Government Marketplace.

Azure Public Cloud Locations (For PAN-OS 7.1, 8.0, and 8.1)		
Americas	Asia Pacific	Europe
• East US	• Australia East	• West Europe
• East US 2	• Australia Southeast	• North Europe
• West US	• East Asia	• UK West
• West US 2	• Japan East	• UK South
• Central US	• Japan West	
• West Central US	• Korea Central	
• North Central US	• Korea South	
• South Central US	• Southeast Asia	
• Canada Central	• China North	
• Canada East	• China East	
• Brazil South		
Azure Government (US) Locations (For PAN-OS 7.1.1, 8.0, and 8.1)		
• Iowa		
• Virginia		
• Texas (PAN-OS 8.1)		
Azure Germany (PAN-OS 8.0)		
Germany Northeast		
Germany Central		

The VM-Series firewall is not supported on the Azure DoD regions.

© 2018 Palo Alto Networks, Inc.

Google Cloud Regions

You can deploy the VM-Series firewall in all Google Cloud Platform regions.

© 2018 Palo Alto Networks, Inc.

AWS CFT Amazon Machine Images (AMI) List

Find the AMI ID that matches the PAN-OS version and licensing option you selected for auto-scaling the VM-Series firewall in AWS.

For the VM-Series Auto Scaling Template for AWS v1.2 you can use the AMIs for 7.1 and 8.0; VM-Series Auto Scaling Template for AWS v2.0 requires the AMI for PAN-OS 8.0.6 or later.

- Images for PAN-OS 8.1
- Images for PAN-OS 8.0
- Images for PAN-OS 7.1.11

Images for PAN-OS 8.1

List of PAN-OS 8.1 AMI IDs for with the different licensing options: BYOL, and PAYG bundle 1 and 2.

- AMI IDs for PAN-OS 8.1.0

AMI IDs for PAN-OS 8.1.0

Use the following AMIs with the VM-Series Auto Scaling template v2.0 and v1.2.

AWS Regions		PAN-OS 8.1.0 AMI IDs that correspond to the licensing options		
		BYOL	PAYG Bundle 2	PAYG Bundle 1
us-east-1	US East (N. Virginia)	ami-a2fa3bdf	ami-bffd3cc2	ami-ce01c0b3
us-east-2	US East (Ohio)	ami-11e1d774	ami-9ef3c5fb	ami-10f3c575
us-west-1	US West (N. California)	ami-a95b4fc9	ami-854551e5	ami-235b4f43
us-west2	US West (Oregon)	ami-d424b5ac	ami-9a29b8e2	ami-8a22b3f2
sa-east-1	South America (SaoPaulo)	ami-9c0154f0	ami-d80653b4	ami-9e0154f2
eu-west-1	EU (Ireland)	ami-62b5fb1b	ami-1fb1ff66	ami-edb0fe94
eu-west-2	EU (London)	ami-876a8de0	ami-c4688fa3	ami-f46a8d93
eu-central-1	EU (Frankfurt)	ami-55bfd73a	ami-1ebdd571	ami-1bbdd574
ca-central-1	Canada (Central)	ami-64038400	ami-57048333	ami-dd0582b9
ap-northeast-1	Asia Pacific (Tokyo)	ami-57662d31	ami-75652e13	ami-39662d5f
ap-northeast-2	Asia Pacific (Seoul)	ami-49bd1127	ami-a8bf13c6	ami-4eb81420

© 2018 Palo Alto Networks, Inc.

AWS Regions		PAN-OS 8.1.0 AMI IDs that correspond to the licensing options		
ap-southeast-1	Asia Pacific (Singapore)	ami-27baeb5b	ami-36bdec4a	ami-55bced29
ap-southeast-2	Asia Pacific (Sydney)	ami-00d61562	ami-add013cf	ami-aed112cc
ap-south-1	Asia Pacific (Mumbai)	ami-e780d988	ami-ee80d981	ami-d385dcbc

Images for PAN-OS 8.0

List of AMI IDs for the different licensing options: bring your own license (BYOL) and pay-as-you-go (PAYG) bundles 1 and 2.

- AMI IDs for PAN-OS 8.0.3
- AMI IDs for PAN-OS 8.0.6
- AMI IDs for PAN-OS 8.0.8
- AMI IDs for PAN-OS 8.0.9

AMI IDs for PAN-OS 8.0.3

Use the following AMIs with the VM-Series Auto Scaling template v1.2.

AWS Regions		PAN-OS 8.0.3 AMI IDs that correspond to the licensing options		
		BYOL	PAYG Bundle 2	PAYG Bundle 1
us-east-1	US East (N. Virginia)	ami-34ca984f	ami-7dcb9906	ami-12c89a69
us-east-2	US East (Ohio)	ami-765e7e13	ami-f1200094	ami-b25e7ed7
us-west-1	US West (N. California)	ami-be88a1de	ami-0f88a16f	ami-8f89a0ef
us-west2	US West (Oregon)	ami-7e8a9107	ami-d28992ab	ami-bb8893c2
sa-east-1	South America (SaoPaulo)	ami-38641354	ami-15651279	ami-1a661176
eu-west-1	EU (Ireland)	ami-1c28c065	ami-86d63eff	ami-82d53dfb
eu-west-2	EU (London)	ami-c76d7ca3	ami-3c170658	ami-3f10015b
eu-central-1	EU (Frankfurt)	ami-44f25e2b	ami-6df35f02	ami-9ff75bf0
ca-central-1	Canada (Central)	ami-0504bb61	ami-0f08b76b	ami-0d0ab569
ap-northeast-1	Asia Pacific (Tokyo)	ami-8705e6e1	ami-ab04e7cd	ami-8105e6e7

© 2018 Palo Alto Networks, Inc.

AWS Regions		PAN-OS 8.0.3 AMI IDs that correspond to the licensing options		
ap-northeast-2	Asia Pacific (Seoul)	ami-48bc6526	ami-6cbd6402	ami-86bb62e8
ap-southeast-1	Asia Pacific (Singapore)	ami-0994066a	ami-17a41074	ami-a79604c4
ap-southeast-2	Asia Pacific (Sydney)	ami-2ed3cc4d	ami-8ed3cced	ami-84d2cde7
ap-south-1	Asia Pacific (Mumbai)	ami-38fd8557	ami-1ffc8470	ami-e0fe868f

AMI IDs for PAN-OS 8.0.6

Use the following AMIs with the VM-Series Auto Scaling template v2.0 and v1.2.

AWS Regions		PAN-OS 8.0.6 AMI IDs that correspond to the licensing options		
		BYOL	PAYG Bundle 2	PAYG Bundle 1
us-east-1	US East (N. Virginia)	ami-493c7233	ami-69a2f513	ami-b6a4f3cc
us-east-2	US East (Ohio)	ami-924269f7	ami-66183303	ami-10193275
us-west-1	US West (N. California)	ami-5d59583d	ami-a29a9bc2	ami-5a9a9b3a
us-west2	US West (Oregon)	ami-7c812804	ami-725ef50a	ami-d258f3aa
sa-east-1	South America (SaoPaulo)	ami-cdd595a1	ami-fb86c597	ami-0787c46b
eu-west-1	EU (Ireland)	ami-d6c44faf	ami-fc4fdb85	ami-844bdffd
eu-west-2	EU (London)	ami-166c7472	ami-e9dcc48d	ami-45d8c021
eu-central-1	EU (Frankfurt)	ami-5b2cba34	ami-a1ad3dce	ami-9151c1fe
ca-central-1	Canada (Central)	ami-3913a95d	ami-50cb4e34	ami-d9cc49bd
ap-northeast-1	Asia Pacific (Tokyo)	ami-4a6eff2c	ami-b5af32d3	ami-11a43977
ap-northeast-2	Asia Pacific (Seoul)	ami-9b2889f5	ami-32f3535c	ami-4bf35325
ap-southeast-1	Asia Pacific (Singapore)	ami-e2cba19e	ami-5ef78022	ami-aff582d3
ap-southeast-2	Asia Pacific (Sydney)	ami-1af80b78	ami-059f6d67	ami-0d9f6d6f

© 2018 Palo Alto Networks, Inc.

AWS Regions		PAN-OS 8.0.6 AMI IDs that correspond to the licensing options		
ap-south-1	Asia Pacific (Mumbai)	ami-dd4316b2	ami-f9005496	ami-e106528e

AMI IDs for PAN-OS 8.0.8

Use the following AMIs with the VM-Series Auto Scaling template v2.0 and v1.2.

AWS Regions		PAN-OS 8.0.8 AMI IDs that correspond to the licensing options		
		BYOL	PAYG Bundle 2	PAYG Bundle 1
us-east-1	US East (N. Virginia)	ami-66c63b1b	ami-ce3cc1b3	ami-b038c5cd
us-east-2	US East (Ohio)	ami-cb8eb9ae	ami-bb8bbcde	ami-0480b761
us-west-1	US West (N. California)	ami-d4f3f9b4	ami-78f1fb18	ami-e0f1fb80
us-west2	US West (Oregon)	ami-4f118437	ami-b21287ca	ami-6d1e8b15
sa-east-1	South America (SaoPaulo)	ami-1ff1ba73	ami-9df0bbf1	ami-1ff7bc73
eu-west-1	EU (Ireland)	ami-a0b5f5d9	ami-cdbfffb4	ami-cbc080b2
eu-west-2	EU (London)	ami-3d93745a	ami-6891760f	ami-3a94735d
eu-central-1	EU (Frankfurt)	ami-aff19ec0	ami-45c8a72a	ami-96caa5f9
ca-central-1	Canada (Central)	ami-d4be39b0	ami-56bd3a32	ami-0abb3c6e
ap-northeast-1	Asia Pacific (Tokyo)	ami-aa662bcc	ami-7a7c311c	ami-e77f3281
ap-northeast-2	Asia Pacific (Seoul)	ami-8109a4ef	ami-8a0fa2e4	ami-610ea30f
ap-southeast-1	Asia Pacific (Singapore)	ami-f61a4e8a	ami-d50054a9	ami-15075369
ap-southeast-2	Asia Pacific (Sydney)	ami-9131f0f3	ami-ae2eefcc	ami-3230f150
ap-south-1	Asia Pacific (Mumbai)	ami-2a471945	ami-8c421ce3	ami-14401e7b

AMI IDs for PAN-OS 8.0.9

Use the following AMIs with the VM-Series Auto Scaling template v2.0 and v1.2.

© 2018 Palo Alto Networks, Inc.

AWS Regions		PAN-OS 8.0.9 AMI IDs that correspond to the licensing options		
		BYOL	PAYG Bundle 2	PAYG Bundle 1
us-east-1	US East (N. Virginia)	ami-49900636	ami-b69305c9	ami-ea950395
us-east-2	US East (Ohio)	ami-8e98a5eb	ami-8998a5ec	ami-3f9aa75a
us-west-1	US West (N. California)	ami-04273e64	ami-26243d46	ami-cb253cab
us-west2	US West (Oregon)	ami-659eec1d	ami-089be970	ami-3398ea4b
sa-east-1	South America (SaoPaulo)	ami-1eb2ee72	ami-70bfe31c	ami-49b4e825
eu-west-1	EU (Ireland)	ami-ef80b096	ami-c69dadbf	ami-e19bab98
eu-west-2	EU (London)	ami-3bcc215c	ami-becf22d9	ami-39cd205e
eu-central-1	EU (Frankfurt)	ami-f3c2ed18	ami-f0c2ed1b	ami-80c3ec6b
ca-central-1	Canada (Central)	ami-9751d1f3	ami-3f51d15b	ami-bd4eced9
ap-northeast-1	Asia Pacific (Tokyo)	ami-621bec1d	ami-691ceb16	ami-491ceb36
ap-northeast-2	Asia Pacific (Seoul)	ami-38b61e56	ami-e7b41c89	ami-3fb61e51
ap-southeast-1	Asia Pacific (Singapore)	ami-39586945	ami-c05968bc	ami-f85a6b84
ap-southeast-2	Asia Pacific (Sydney)	ami-788a5b1a	ami-958b5af7	ami-ce8a5bac
ap-south-1	Asia Pacific (Mumbai)	ami-d06e43bf	ami-ca6b46a5	ami-5169443e

Images for PAN-OS 7.1.11

List of AMI IDs for PAN-OS 7.1.11 with the different licensing options: BYOL, and PAYG bundle 1 and 2.

AWS Regions		PAN-OS 7.1.11 AMI IDs that correspond to the licensing options		
		BYOL	PAYG Bundle 2	PAYG Bundle 1

© 2018 Palo Alto Networks, Inc.

AWS Regions		PAN-OS 7.1.11 AMI IDs that correspond to the licensing options		
us-east-1	US East (N. Virginia)	ami-83f1a3f8	ami-6ff0a214	ami-19f3a162
us-east-2	US East (Ohio)	ami-d15e7eb4	ami-a72101c2	ami-ba5e7edf
us-west-1	US West (N. California)	ami-0189a061	ami-0089a060	ami-4f88a12f
us-west2	US West (Oregon)	ami-0a829973	ami-d28398ab	ami-788e9501
sa-east-1	South America (Sao Paulo)	ami-3d631451	ami-c26413ae	ami-3c631450
eu-west-1	EU (Ireland)	ami-6a21c913	ami-1c21c965	ami-2920c850
eu-west-2	EU (London)	ami-176e7f73	ami-166e7f72	ami-b31302d7
eu-central-1	EU (Frankfurt)	ami-80f35fef	ami-1ecc6071	ami-4acc6025
ca-central-1	Canada (Central)	ami-1c02bd78	ami-440eb120	ami-1303bc77
ap-northeast-1	Asia Pacific (Tokyo)	ami-750dee13	ami-220eed44	ami-710cef17
ap-northeast-2	Asia Pacific (Seoul)	ami-06ba6368	ami-74bc651a	ami-96ba63f8
ap-southeast-1	Asia Pacific (Singapore)	ami-cc9a08af	ami-e2950781	ami-ae9a08cd
ap-southeast-2	Asia Pacific (Sydney)	ami-10d3cc73	ami-13d3cc70	ami-42d3cc21
ap-south-1	Asia Pacific (Mumbai)	ami-d4fe86bb	ami-2dfc8442	ami-a5fe86ca

Panorama

This section includes information on Panorama and compatible versions for devices that Panorama can manage, and on plugins available for Panorama.

- Plugins
- Panorama Hypervisor Support

© 2018 Palo Alto Networks, Inc.

Plugins

Panorama 8.0 introduces support for the extensible plugin architecture. The following tables describe the features and functionality introduced with the Panorama plugins.

- VMware NSX
- Cloud Services
- Interconnect on page 34
- Public Cloud-AWS, Azure, GCP on page 34

VMware NSX

The following table shows the features introduced in each version of the VM-Series firewall VMware NSX plugin. For additional information about each plugin, see the release notes on the Customer Support Portal.

Plugin Version	Minimum Panorama Version	New Features or Changes
1.0.0	8.0.0	First plugin introduced to support the integration with VMware NSX.
1.0.1	8.0.0	Adds support for NSX Manager 6.3 with vSphere 6.5a.
1.0.2	8.0.0	Introduces fixes for known issues. No new features.
1.0.3	8.0.0	Introduces fixes for known issues. No new features.
2.0.0	8.0.5	Introduces support for multiple NSX Managers on a single Panorama.
2.0.1	8.0.5	Introduces fixes for known issues. Minimum required plugin version for Panorama 8.1.
2.0.2	8.0.5	Introduces fixes for known issues.
2.0.3	8.0.5	Introduces fixes for known issues.

Cloud Services

The Cloud Services plugin supports the GlobalProtect cloud service and the Logging service:

© 2018 Palo Alto Networks, Inc.

Panorama Minimum Version	Cloud Services Plugin Minimum Version
8.0.5	1.2.0
8.1.0	1.2.0
8.1.1	1.2.0

For additional information about this plugin, see the Release Notes.

Interconnect

The following table shows the features introduced in each version of the Panorama Interconnect plugin.

Plugin Version	Minimum Panorama Version	New Features or Changes
1.0.0	8.1.3	First plugin introduced to support a two-tier Panorama deployment for a horizontal scale-out architecture.

Public Cloud-AWS, Azure, GCP

The following table shows the features introduced in each version of the Panorama plugin for Amazon Web Services, Microsoft Azure, and Google Cloud Platform. The plugins use device groups and templates on Panorama to push the configuration to the managed firewalls.

Public Cloud Platform	Plugin Version	Minimum Panorama Version	Features
AWS	1.0.0	8.1.3	Enables support for VM Monitoringfrom Panorama. Configure the Panorama plugin for AWS to monitor the virtual machine inventory within your AWS VPCs so that you can consistently enforce Securitypolicy that automatically adapts to changes within your AWS deployment.
Azure	1.0.0	8.1.3	Enables support for VM Monitoring from Panorama. Configure the Panorama plugin for Azure to monitor the virtual machine inventory within your Azure subscriptions so that you can consistently enforce Securitypolicy that automatically adapts to changes within your Azure deployment.

© 2018 Palo Alto Networks, Inc.

Public Cloud Platform	Plugin Version	Minimum Panorama Version	Features
GCP	1.0.0	8.1.3	Enables you to secure Kubernetes services in a Google Kubernetes Engine (GKE) cluster. Configure the Panorama plugin for GCP to connect to your GKE cluster and learn about the services that are exposed to the internet.

© 2018 Palo Alto Networks, Inc.

Panorama Hypervisor Support

Before you install Panorama on a VMware ESXi or Hyper-V server, verify that the hypervisor meets the minimum version requirements to deploy Panorama.

Panorama Version	VMware ESXi Compatibility	Hyper-V Compatibility	Public Cloud Compatibility
PAN-OS 6.1	64-bit kernel-based VMware ESX(i) 5.1 or 5.5.	Not Available	Not Available
PAN-OS 7.0	64-bit kernel-based VMware ESXi 5.1, 5.5, or 6.0. The minimum supported version of the virtual hardware family type (also known as the VMware virtual hardware version) on the ESXi server is vmx-09.	Not Available	Not Available
PAN-OS 7.1	64-bit kernel-based VMware ESXi 5.1, 5.5, 6.0, or 6.5. The minimum supported version of the virtual hardware family type (also known as the VMware virtual hardware version) on the ESXi server is vmx-09. ESXi 5.5 and later versions supports a virtual disk of up to8TB. ESXi versions support one disk of up to 2TB.	Not Available	Not Available
PAN-OS 8.0	64-bit kernel-based VMware ESXi 5.1, 5.5, 6.0, or 6.5. The supported version of the virtual hardware family type (also known as the VMware virtual hardware version) on the ESXi server is vmx-09.	Not Available	Not Available

© 2018 Palo Alto Networks, Inc.

Panorama Version	VMware ESXi Compatibility	Hyper-V Compatibility	Public Cloud Compatibility
	ESXi 5.5 and later versions supports one disk of up to 8TB. Earlier ESXi versions support one disk of up to 2TB.		
PAN-OS 8.1	64-bit kernel-based VMware ESXi 5.1, 5.5, 6.0, or 6.5. The supported version of the virtual hardware family type (also known as the VMware virtual hardware version) on the ESXi server is vmx-09. ESXi 5.5 and later versions supports one disk of up to 8TB. Earlier ESXi versions support one disk of up to 2TB.	(8.1.3) Windows Server 2016 with Hyper-V role or Hyper-V 2016	• Amazon AWS • Microsoft Azure • (8.1.1) Google Cloud Platform • (8.1.3) Amazon AWS GovCloud

© 2018 Palo Alto Networks, Inc.

MFA Vendor Support

© 2018 Palo Alto Networks, Inc.

MFA Vendor Support

For the following authentication use cases, Palo Alto Networks® appliances can integrate with multi-factor authentication (MFA) vendors using RADIUS and, in PAN-OS® 8.0, using SAML:

- Remote user authentication through the GlobalProtect™ portals and gateways.
- Administrator authentication in the PAN-OS and Panorama™ web interface.
- Authentication through the Captive Portal policy (PAN-OS 7.1 and earlier releases) and Authentication policy (PAN-OS 8.0 and later releases).

Additionally, firewalls running a PAN-OS 8.0 or later release can also integrate with the following MFA vendors using the API to enforce MFA through Authentication policy.

Palo Alto Networks provides support for MFA vendors through Applications content updates. This means that if you use Panorama to push device group configurations to firewalls, you must install the same Applications release version on the firewalls as on Panorama to avoid mismatches in vendor support.

MFA Vendor	Minimum Content Release Version
Duo v2	655
Okta Adaptive	655
PingID	655
RSA SecurID Access	752

© 2018 Palo Alto Networks, Inc.

Supported Cipher Suites

Use this table in the Palo Alto Networks® Compatibility Matrix to determine support for cipher suites according to function and PAN-OS® release.

- Cipher Suites Supported in PAN-OS 8.1
- Cipher Suites Supported in PAN-OS 8.0
- Cipher Suites Supported in PAN-OS 7.1

© 2018 Palo Alto Networks, Inc.

Cipher Suites Supported in PAN-OS 8.1

The following topics list cipher suites that are supported on firewalls running a PAN-OS 8.1 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.1 Cipher Suites Supported in FIPS-CC Mode.

The ciphers supported in normal operation mode are grouped according to feature or functionality in the following sections:

- PAN-OS 8.1 GlobalProtect Cipher Suites
- PAN-OS 8.1 HA1 SSH Cipher Suites on page 51
- PAN-OS 8.1 IPSec Cipher Suites
- PAN-OS 8.1 IKE and Web Certificate Cipher Suites
- PAN-OS 8.1 Decryption Cipher Suites
- PAN-OS 8.1 Administrative Session Cipher Suites
- PAN-OS 8.1 PAN-OS-to-Panorama Connection Cipher Suites

PAN-OS 8.1 GlobalProtect Cipher Suites

The following table lists cipher suites for GlobalProtect™ supported on firewalls running a PAN-OS® 8.1 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.1 Cipher Suites Supported in FIPS-CC Mode*.*

- GlobalProtect App/Agent—SSL tunnels and SSL connections to gateway and portal
- GlobalProtect App/Agent—IPSec mode
- GlobalProtect Portal—Browser Access

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
GlobalProtect App/Agent—SSL tunnels and SSL connections to gateway and portal	• TLSv1.0, TLSv1.1, and TLSv1.2 cipher suites • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-3DES-SHA • RSA-AES-128-SHA • RSA-AES-256-SHA • RSA-AES-128-SHA-256 • RSA-AES-256-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-SEED-SHA • DHE-RSA-AES-128-SHA • DHE-RSA-AES-256-SHA • DHE-RSA-AES-128-GCM-SHA-256 • DHE-RSA-AES-256-GCM-SHA-384 • EDH-RSA-3DES-SHA

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
	• ECDHE-RSA-AES-128-SHA • ECDHE-RSA-AES-256-SHA • ECDHE-RSA-AES-128-GCM-SHA-256 • ECDHE-RSA-AES-128-GCM-SHA-384 • ECDHE-ECDSA-AES-128-SHA • ECDHE-ECDSA-AES-256-SHA • ECDHE-ECDSA-AES-128-GCM-SHA-256 • ECDHE-ECDSA-AES-256-GCM-SHA-384
GlobalProtect App/Agent—IPSec mode (Keys transported through SSL session with gateway)	• AES-128-CBC-HMAC-SHA • AES-128-GCM-HMAC-SHA • AES-256-GCM-HMAC-SHA
GlobalProtect Portal—Browser Access	• SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 cipher suites • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-3DES-SHA • RSA-AES-128-SHA • RSA-AES-256-SHA • RSA-AES-128-SHA-256 • RSA-AES-256-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-AES-256-SHA • DHE-RSA-AES-128-SHA • DHE-RSA-AES-128-GCM-SHA-256 • DHE-RSA-AES-256-GCM-SHA-384 • EDH-RSA-3DES-SHA • ECDHE-ECDSA-AES-128-SHA • ECDHE-ECDSA-AES-256-SHA • ECDHE-ECDSA-AES-128-GCM-SHA-256 • ECDHE-ECDSA-AES-256-GCM-SHA-384

PAN-OS 8.1 IPSec Cipher Suites

The following table lists the cipher suites for IPSec that are supported on firewalls running a PAN-OS® 8.1 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.1 Cipher Suites Supported in FIPS-CC Mode.

- IPSec—Encryption
- IPSec—Message Authentication
- IPSec—Key Exchange

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
IPSec—Encryption	• NULL • DES • 3DES • AES-128-CBC • AES-192-CBC • AES-256-CBC • AES-128-CCM • AES-128-GCM • AES-256-GCM
IPSec—Message Authentication	• NONE • HMAC-MD5 • HMAC-SHA • HMAC-SHA-256 • HMAC-SHA-384 • HMAC-SHA-512
IPSec—Key Exchange	Diffie-Hellman groups with or without perfect forward secrecy (PFS): • No PFS—This option specifies that the firewall reuses the same key for IKE phase 1 and phase 2 instead of renewing the key for phase 2. • Group 1 (768-bit keys) with PFS enabled • Group 2 (1024-bit keys) with PFS enabled • Group 5 (1536-bit keys) with PFS enabled • Group 14 (2048-bit keys) with PFS enabled • Group 19 (256-bit elliptic curve group) with PFS enabled • Group 20 (384-bit elliptic curve group) with PFS enabled

PAN-OS 8.1 IKE and Web Certificate Cipher Suites

The following table lists cipher suites for Internet Key Exchange (IKE) and PAN-OS® web certificates that are supported on firewalls running a PAN-OS 8.1 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.1 Cipher Suites Supported in FIPS-CC Mode.

- IKE Certificate Support
- IKE—Encryption
- IKE—Message Authentication
- IKE—Key Exchange
- PAN-OS Web Certificates

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
IKE Certificate Support	• RSA • Keys—512-, 1024-, 2048-, and 3072-bit keys

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases

	• Digital signature algorithms—SHA, SHA-256, SHA-384, or SHA-512
	• ECDSA
	  • Keys—256- and 384-bit keys
	  • Digital signature algorithms—SHA-256, SHA-384, or SHA-512

IKE—Encryption	• DES • 3DES • AES-128-CBC • AES-192-CBC • AES-256-CBC
IKE—Message Authentication	• HMAC-MD5 • HMAC-SHA • HMAC-SHA-256 • HMAC-SHA-384 • HMAC-SHA-512
IKE—Key Exchange	Diffie-Hellman groups • Group 1 (768-bit keys) • Group 2 (1024-bit keys) • Group 5 (1536-bit keys) • Group 14 (2048-bit keys) • Group 19 (256-bit elliptic curve group) • Group 20 (384-bit elliptic curve group)

PAN-OS Web Certificates	• RSA
	  • Keys—512-, 1024-, 2048-, 3072- and 4096-bit keys
	  • Digital signature algorithms—SHA, SHA-256, SHA-384, or SHA-512
	• ECDSA
	  • Keys—256- and 384-bit keys
	  • Digital signature algorithms—SHA-256, SHA-384, or SHA-512

PAN-OS 8.1 Decryption Cipher Suites

The following table lists cipher suites for decryption that are supported on firewalls running a PAN-OS® 8.1 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.1 Cipher Suites Supported in FIPS-CC Mode.

- SSH Decryption (SSHv2 only)—Encryption
- SSH Decryption (SSHv2 only)—Message Authentication
- SSL/TLS Decryption
- SSL/TLS Decryption—NIST-approved Elliptical Curves
- SSL/TLS Decryption—Perfect Forward Secrecy (PFS) Ciphers

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
SSH Decryption (SSHv2 only)—Encryption	• AES-128-CBC • AES-192-CBC • AES-256-CBC • AES-128-CTR • AES-192-CTR • AES-256-CTR
SSH Decryption (SSHv2 only)—Message Authentication	• HMAC-RIPEMD • HMAC-MD5-96 • HMAC-MD5 • HMAC-SHA-96 • HMAC-RIPEMD-160 • HMAC-SHA
SSL/TLS Decryption	• SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 cipher suites • RSA 512-, 1024-, 2048-, 3072-, 4096-, and 8192-bit keys  *The firewall can authenticate certificates up to 8192-bit RSA keys from the destination server, however the firewall generated certificate to the client supports only up to 2048-bit RSA keys.* • RSA-RC4-128-MD5 • RSA-RC4-128-SHA • RSA-3DES-EDE-CBC-SHA • RSA-AES-128-CBC-SHA • RSA-AES-256-CBC-SHA • RSA-AES-128-CBC-SHA-256 • RSA-AES-256-CBC-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384
SSL/TLS Decryption—NIST-approved Elliptical Curves	• P-192 (secp192r1) • P-224 (secp224r1) • P-256 (secp256r1) • P-384 (secp384r1) • P-521 (secp521r1)
SSL/TLS Decryption—Perfect Forward Secrecy (PFS) Ciphers  *If you use the DHE or ECDHE key exchange algorithms to enable PFS, you cannot use*	• DHE-RSA-3DES-EDE-CBC-SHA • DHE-RSA-AES-128-CBC-SHA • DHE-RSA-AES-256-CBC-SHA • DHE-RSA-AES-128-CBC-SHA-256 • DHE-RSA-AES-256-CBC-SHA-256 • DHE-RSA-AES-128-GCM-SHA-256 • DHE-RSA-AES-256-GCM-SHA-384 • ECDHE-RSA-AES-128-CBC-SHA • ECDHE-RSA-AES-256-CBC-SHA • ECDHE-RSA-AES-128-CBC-SHA-256 • ECDHE-RSA-AES-256-CBC-SHA-384

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
a hardware security module (HSM) to store the private keys used for SSL Inbound Inspection.	• ECDHE-RSA-AES-128-GCM-SHA-256 • ECDHE-RSA-AES-256-GCM-SHA-384 • ECDHE-ECDSA-AES-128-CBC-SHA • ECDHE-ECDSA-AES-256-CBC-SHA • ECDHE-ECDSA-AES-128-CBC-SHA-256 • ECDHE-ECDSA-AES-256-CBC-SHA-384 • ECDHE-ECDSA-AES-128-GCM-SHA-256 • ECDHE-ECDSA-AES-256-GCM-SHA-384

PAN-OS 8.1 Administrative Session Cipher Suites

The following table lists the cipher suites for administrative sessions that are supported on firewalls running a PAN-OS® 8.1 release in normal (non-FIPS-CC) operational mode.



If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.1 Cipher Suites Supported in FIPS-CC Mode*.*

- Administrative Sessions to Web Interface
- Administrative Sessions to CLI (SSH)—Encryption
- Administrative Sessions to CLI (SSH)—Message Authentication
- Administrative Sessions to CLI (SSH)—Server Host Key Types
- Administrative Sessions to CLI (SSH)—Key Exchange Algorithms

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
Administrative Sessions to Web Interface	• TLSv1.1 and TLSv1.2 cipher suites • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-3DES-SHA • RSA-AES-128-SHA • RSA-AES-256-SHA • RSA-AES-256-CBC-SHA • RSA-AES-128-CBC-SHA-256 • RSA-AES-256-CBC-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-3DES-SHA • DHE-RSA-AES-128-GCM-SHA-256 • DHE-RSA-AES-256-GCM-SHA-384 • ECDHE-RSA-AES-128-GCM-SHA-256 • ECDHE-RSA-AES-256-GCM-SHA-384 • ECDHE-ECDSA-AES-128-SHA • ECDHE-ECDSA-AES-256-SHA • ECDHE-ECDSA-AES-128-GCM-SHA-256

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
	• ECDHE-ECDSA-AES-256-GCM-SHA-384
Administrative Sessions to CLI (SSH)—Encryption	• 3DES • ARCFOUR128 • ARCFOUR256 • BLOWFISH • CAST128 • AES-128-CBC • AES-192-CBC • AES-256-CBC • AES-128-CTR • AES-192-CTR • AES-256-CTR • AES-128-GCM • AES-256-GCM
Administrative Sessions to CLI (SSH)—Message Authentication	• UMAC-64 • HMAC-MD5-96 • HMAC-MD5 • HMAC-SHA-96 • HMAC-RIPEMD-160 • HMAC-SHA • HMAC-SHA-256 • HMAC-SHA-512
Administrative Sessions to CLI (SSH)—Server Host Key Types	• RSA keys—2048-, 3072-, and 4096-bit keys • ECDSA keys—256-, 384-, and 521-bit keys
Administrative Sessions to CLI (SSH)—Key Exchange Algorithms	• diffie-hellman-group1-sha1 • diffie-hellman-group14-sha1 • diffie-hellman-group-exchange-sha1 • diffie-hellman-group-exchange-sha256 • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521

PAN-OS 8.1 HA1 SSH Cipher Suites

The following table lists the cipher suites for HA1 control connections using SSH that are supported on firewalls running a PAN-OS 8.1 release in normal (non-FIPS-CC) or FIPS-CC operational mode.

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
HA1 SSH	• AES 128-bit cipher with Cipher Block Chaining • AES 128-bit cipher with Counter Mode • AES 128-bit cipher with GCM (Galois/Counter Mode)

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
	• AES 192-bit cipher with Cipher Block Chaining • AES 192-bit cipher with Counter Mode • AES 256-bit cipher with Cipher Block Chaining • AES 256-bit cipher with Counter Mode • AES 256-bit cipher with GCM

PAN-OS 8.1 PAN-OS-to-Panorama Connection Cipher Suites

The following table lists the cipher suites for PAN-OS®-to-Panorama™ connections that are supported on firewalls running a PAN-OS 8.1 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.1 Cipher Suites Supported in FIPS-CC Mode.

Feature or Function	Ciphers Supported in PAN-OS 8.1 Releases
PAN-OS to Panorama Connection	• RSA-RC4-128-SHA • RSA-3DES-SHA • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-AES-128-SHA • RSA-AES-128-SHA-256 • RSA-AES-256-SHA • RSA-AES-256-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-AES-128-SHA • DHE-RSA-AES-256-SHA

PAN-OS 8.1 Cipher Suites Supported in FIPS-CC Mode

The following table lists cipher suites that are supported on firewalls running a PAN-OS® 8.1 release in FIPS-CC mode.

If your firewall is running in normal (non-FIPS-CC) operational mode, see Cipher Suites Supported in PAN-OS 8.1

Functions	Standards	Certificates
Asymmetric key generation		
FFC key pair generation (key size 2048 bits)	FIPS PUB 186-4	**Appliances:** DSA #TBD **VMs:**

© 2018 Palo Alto Networks, Inc.

Functions	Standards	Certificates
		DSA #TBD
ECC key pair generation (NIST curves P-256, P-384)	FIPS PUB 186-4	**Appliances:** ECDSA #TBD **VMs:** ECDSA #TBD
RSA key generation (key size 2048 bits)	FIPS PUB 186-4	**Appliances:** RSA #TBD **VMs:** RSA #TBD
Cryptographic Key Generation (for IKE Peer Authentication)		
RSA key generation (key size 2048 bits)	FIPS PUB 186-4	**Appliances:** RSA #TBD **VMs:** RSA #TBD
ECDSAkey pair generation (NIST curves P-256, P-384)	FIPS PUB 186-4	**Appliances:** ECDSA #TBD **VMs:** ECDSA #TBD
Cryptographic Key Establishment		
ECDSA based key establishment	NIST SP 800-56A Revision 2	**Appliances:** CVL #TBD **VMs:** CVL #TBD
FFC based key establishment	NIST SP 800-56A Revision 2	**Appliances:** CVL #TBD **VMs:** CVL #TBD
AES Data Encryption/Decryption		
• AES 128 and AES 256 • CBC 128 and CBC 256 • GCM 128 and GCM 256	• AES as specified in ISO 18033-3 • CBC as specified in ISO 10116	**Appliances:** AES #TBD **VMs:**

© 2018 Palo Alto Networks, Inc.

Functions	Standards	Certificates
	• GCM as specified in ISO 19772	AES #TBD
Signature Generation and Verification		
RSA Digital Signature Algorithm (rDSA) (modulus 2048)	FIPS PUB 186-4, “Digital Signature Standard (DSS)”, Section 5.5, using PKCS #1 v2.1 Signature Schemes RSASSA-PSS and/or RSASSAPKCS1v1_5; ISO/IEC 9796-2, Digital signature scheme 2 or Digital Signature scheme 3	**Appliances:** RSA #TBD **VMs:** RSA #TBD
ECDSA (NIST curves P-256, P-384, and P-521)	FIPS PUB 186-4, “Digital Signature Standard (DSS)”, Section 6 and Appendix D, Implementing "NIST curves" P-256, P-384, ISO/IEC 14888-3, Section 6.4	**Appliances:** RSA #TBD **VMs:** RSA #TBD
Cryptographic hashing		
SHA-1, SHA-256, SHA-384 and SHA-512 (digest sizes 160, 256, 384 and 512 bits)	ISO/IEC 10118-3:2004	**Appliances:** SHS #TBD **VMs:** SHS #TBD
Keyed-hash message authentication		
• HMAC-SHA-1 • MAC-SHA-256 • HMAC-SHA-384 • HMAC-SHA-512	ISO/IEC 9797-2:2011	**Appliances:** HMAC #TBD **VMs:** HMAC #TBD
Random bit generation		
CTR_DRBG (AES) from a hardware based noise source with one independent software-based noise source of 256 bits of non-determinism	ISO/IEC 18031:2011	**Appliances:** DRBG #TBD **VMs:** DRBG #TBD

© 2018 Palo Alto Networks, Inc.

Cipher Suites Supported in PAN-OS 8.0

The following table lists cipher suites that are supported on firewalls running a PAN-OS 8.0 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.0 Cipher Suites Supported in FIPS-CC Mode.

The ciphers supported in normal operation mode are grouped according to feature or functionality in the following sections:

- PAN-OS 8.0 GlobalProtect Cipher Suites
- PAN-OS 8.0 IPSec Cipher Suites
- PAN-OS 8.0 IKE and Web Certificate Cipher Suites
- PAN-OS 8.0 Decryption Cipher Suites
- PAN-OS 8.0 Administrative Session Cipher Suites
- PAN-OS 8.0 PAN-OS-to-Panorama Connection Cipher Suites

PAN-OS 8.0 GlobalProtect Cipher Suites

The following table lists cipher suites for GlobalProtect™ supported on firewalls running a PAN-OS® 8.0 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.0 Cipher Suites Supported in FIPS-CC Mode*.*

- GlobalProtect App/Agent—SSL tunnels and SSL connections to gateway and portal
- GlobalProtect App/Agent—IPSec mode
- GlobalProtect Portal—Browser Access

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases
GlobalProtect App/Agent—SSL tunnels and SSL connections to gateway and portal	• TLSv1.0, TLSv1.1, and TLSv1.2 cipher suites • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-3DES-SHA • RSA-AES-128-SHA • RSA-AES-256-SHA • RSA-AES-128-SHA-256 • RSA-AES-256-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-SEED-SHA • DHE-RSA-AES-128-SHA • DHE-RSA-AES-256-SHA • DHE-RSA-AES-128-GCM-SHA-256 • DHE-RSA-AES-256-GCM-SHA-384 • EDH-RSA-3DES-SHA • ECDHE-RSA-AES-128-SHA

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases
	• ECDHE-RSA-AES-256-SHA • ECDHE-RSA-AES-128-GCM-SHA-256 • ECDHE-RSA-AES-128-GCM-SHA-384 • ECDHE-ECDSA-AES-128-SHA • ECDHE-ECDSA-AES-256-SHA • ECDHE-ECDSA-AES-128-GCM-SHA-256 • ECDHE-ECDSA-AES-256-GCM-SHA-384
GlobalProtect App/Agent—IPSec mode (Keys transported through SSL session with gateway)	• AES-128-CBC-HMAC-SHA • AES-128-GCM-HMAC-SHA • AES-256-GCM-HMAC-SHA
GlobalProtect Portal—Browser Access	• SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 cipher suites • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-3DES-SHA • RSA-AES-128-SHA • RSA-AES-256-SHA • RSA-AES-128-SHA-256 • RSA-AES-256-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-AES-256-SHA • DHE-RSA-AES-128-SHA • DHE-RSA-AES-128-GCM-SHA-256 • DHE-RSA-AES-256-GCM-SHA-384 • EDH-RSA-3DES-SHA • ECDHE-ECDSA-AES-128-SHA • ECDHE-ECDSA-AES-256-SHA • ECDHE-ECDSA-AES-128-GCM-SHA-256 • ECDHE-ECDSA-AES-256-GCM-SHA-384

PAN-OS 8.0 IPSec Cipher Suites

The following table lists the cipher suites for IPSec that are supported on firewalls running a PAN-OS® 8.0 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.0 Cipher Suites Supported in FIPS-CC Mode*.*

- IPSec—Encryption
- IPSec—Message Authentication
- IPSec—Key Exchange

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases
IPSec—Encryption	• NULL • DES • 3DES • AES-128-CBC • AES-192-CBC • AES-256-CBC • AES-128-CCM • AES-128-GCM • AES-256-GCM
IPSec—Message Authentication	• NONE • HMAC-MD5 • HMAC-SHA • HMAC-SHA-256 • HMAC-SHA-384 • HMAC-SHA-512
IPSec—Key Exchange	Diffie-Hellman groups with or without perfect forward secrecy (PFS): • No PFS—This option specifies that the firewall reuses the same key for IKE phase 1 and phase 2 instead of renewing the key for phase 2. • Group 1 (768-bit keys) with PFS enabled • Group 2 (1024-bit keys) with PFS enabled • Group 5 (1536-bit keys) with PFS enabled • Group 14 (2048-bit keys) with PFS enabled • Group 19 (256-bit elliptic curve group) with PFS enabled • Group 20 (384-bit elliptic curve group) with PFS enabled

PAN-OS 8.0 IKE and Web Certificate Cipher Suites

The following table lists cipher suites for Internet Key Exchange (IKE) and PAN-OS® web certificates that are supported on firewalls running a PAN-OS 8.0 release in normal (non-FIPS-CC) operational mode.



If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.0 Cipher Suites Supported in FIPS-CC Mode.

- IKE Certificate Support
- IKE—Encryption
- IKE—Message Authentication
- IKE—Key Exchange
- PAN-OS Web Certificates

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases
IKE Certificate Support	• RSA • Keys—512-, 1024-, 2048-, and 3072-bit keys

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases

	• Digital signature algorithms—SHA, SHA-256, SHA-384, or SHA-512
	• ECDSA
	  • Keys—256- and 384-bit keys
	  • Digital signature algorithms—SHA-256, SHA-384, or SHA-512

IKE—Encryption	• DES • 3DES • AES-128-CBC • AES-192-CBC • AES-256-CBC
IKE—Message Authentication	• HMAC-MD5 • HMAC-SHA • HMAC-SHA-256 • HMAC-SHA-384 • HMAC-SHA-512
IKE—Key Exchange	Diffie-Hellman groups • Group 1 (768-bit keys) • Group 2 (1024-bit keys) • Group 5 (1536-bit keys) • Group 14 (2048-bit keys) • Group 19 (256-bit elliptic curve group) • Group 20 (384-bit elliptic curve group)

PAN-OS Web Certificates	• RSA
	  • Keys—512-, 1024-, 2048-, 3072- and 4096-bit keys
	  • Digital signature algorithms—SHA, SHA-256, SHA-384, or SHA-512
	• ECDSA
	  • Keys—256- and 384-bit keys
	  • Digital signature algorithms—SHA-256, SHA-384, or SHA-512

PAN-OS 8.0 Decryption Cipher Suites

The following table lists cipher suites for decryption that are supported on firewalls running a PAN-OS® 8.0 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.0 Cipher Suites Supported in FIPS-CC Mode.

- SSH Decryption (SSHv2 only)—Encryption
- SSH Decryption (SSHv2 only)—Message Authentication
- SSL/TLS Decryption
- SSL/TLS Decryption—NIST-approved Elliptical Curves
- SSL/TLS Decryption—Perfect Forward Secrecy (PFS) Ciphers

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases
SSH Decryption (SSHv2 only)—Encryption	• AES-128-CBC • AES-192-CBC • AES-256-CBC • AES-128-CTR • AES-192-CTR • AES-256-CTR
SSH Decryption (SSHv2 only)—Message Authentication	• HMAC-RIPEMD • HMAC-MD5-96 • HMAC-MD5 • HMAC-SHA-96 • HMAC-RIPEMD-160 • HMAC-SHA
SSL/TLS Decryption	• SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 cipher suites • RSA 512-, 1024-, 2048-, 3072-, 4096-, and 8192-bit keys  *The firewall can authenticate certificates up to 8192-bit RSA keys from the destination server, however the firewall generated certificate to the client supports only up to 2048-bit RSA keys.* • RSA-RC4-128-MD5 • RSA-RC4-128-SHA • RSA-3DES-EDE-CBC-SHA • RSA-AES-128-CBC-SHA • RSA-AES-256-CBC-SHA • RSA-AES-128-CBC-SHA-256 • RSA-AES-256-CBC-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384
SSL/TLS Decryption—NIST-approved Elliptical Curves	• P-192 (secp192r1) • P-224 (secp224r1) • P-256 (secp256r1) • P-384 (secp384r1) • P-521 (secp521r1)
SSL/TLS Decryption—Perfect Forward Secrecy (PFS) Ciphers  *If you use the DHE or ECDHE key exchange algorithms to enable PFS, you cannot use*	• DHE-RSA-3DES-EDE-CBC-SHA • DHE-RSA-AES-128-CBC-SHA • DHE-RSA-AES-256-CBC-SHA • DHE-RSA-AES-128-GCM-SHA-256 • DHE-RSA-AES-256-GCM-SHA-384 • ECDHE-RSA-AES-128-CBC-SHA • ECDHE-RSA-AES-256-CBC-SHA • ECDHE-RSA-AES-128-CBC-SHA-256 • ECDHE-RSA-AES-128-GCM-SHA-256 • ECDHE-RSA-AES-256-GCM-SHA-384 • ECDHE-ECDSA-AES-128-CBC-SHA

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases
a hardware security module (HSM) to store the private keys used for SSL Inbound Inspection.	• ECDHE-ECDSA-AES-256-CBC-SHA • ECDHE-ECDSA-AES-128-CBC-SHA-256 • ECDHE-ECDSA-AES-256-CBC-SHA-384 • ECDHE-ECDSA-AES-128-GCM-SHA-256 • ECDHE-ECDSA-AES-256-GCM-SHA-384 Additional ciphers supported in PAN-OS 8.0.6 and later releases: • DHE-RSA-AES-128-CBC-SHA-256 • DHE-RSA-AES-256-CBC-SHA-256 • ECDHE-RSA-AES-128-CBC-SHA-256 • ECDHE-RSA-AES-256-CBC-SHA-384

PAN-OS 8.0 Administrative Session Cipher Suites

The following table lists the cipher suites for administrative sessions that are supported on firewalls running a PAN-OS® 8.0 release in normal (non-FIPS-CC) operational mode.

If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.0 Cipher Suites Supported in FIPS-CC Mode.

- Administrative Sessions to Web Interface
- Administrative Sessions to CLI (SSH)—Encryption
- Administrative Sessions to CLI (SSH)—Message Authentication
- Administrative Sessions to CLI (SSH)—Server Host Key Types
- Administrative Sessions to CLI (SSH)—Key Exchange Algorithms

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases
Administrative Sessions to Web Interface	• TLSv1.1 and TLSv1.2 cipher suites • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-3DES-SHA • RSA-AES-128-SHA • RSA-AES-256-SHA • RSA-AES-256-CBC-SHA • RSA-AES-128-CBC-SHA-256 • RSA-AES-256-CBC-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-3DES-SHA • DHE-RSA-AES-128-GCM-SHA-256 • DHE-RSA-AES-256-GCM-SHA-384 • ECDHE-RSA-AES-128-GCM-SHA-256 • ECDHE-RSA-AES-256-GCM-SHA-384 • ECDHE-ECDSA-AES-128-SHA • ECDHE-ECDSA-AES-256-SHA

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases
	• ECDHE-ECDSA-AES-128-GCM-SHA-256 • ECDHE-ECDSA-AES-256-GCM-SHA-384
Administrative Sessions to CLI (SSH)—Encryption	• 3DES • ARCFOUR128 • ARCFOUR256 • BLOWFISH • CAST128 • AES-128-CBC • AES-192-CBC • AES-256-CBC • AES-128-CTR • AES-192-CTR • AES-256-CTR • AES-128-GCM • AES-256-GCM
Administrative Sessions to CLI (SSH)—Message Authentication	• UMAC-64 • HMAC-MD5-96 • HMAC-MD5 • HMAC-SHA-96 • HMAC-RIPEMD-160 • HMAC-SHA • HMAC-SHA-256 • HMAC-SHA-512
Administrative Sessions to CLI (SSH)—Server Host Key Types	• RSA keys—2048-, 3072-, and 4096-bit keys • ECDSA keys—256-, 384-, and 521-bit keys
Administrative Sessions to CLI (SSH)—Key Exchange Algorithms	• diffie-hellman-group1-sha1 • diffie-hellman-group14-sha1 • diffie-hellman-group-exchange-sha1 • diffie-hellman-group-exchange-sha256 • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521

PAN-OS 8.0 PAN-OS-to-Panorama Connection Cipher Suites

The following table lists the cipher suites for PAN-OS®-to-Panorama™ connections that are supported on firewalls running a PAN-OS 8.0 release in normal (non-FIPS-CC) operational mode.



If your firewall is running in FIPS-CC mode, see the list of PAN-OS 8.0 Cipher Suites Supported in FIPS-CC Mode.

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 8.0 Releases
PAN-OS to Panorama Connection	• RSA-RC4-128-SHA • RSA-3DES-SHA • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-AES-128-SHA • RSA-AES-128-SHA-256 • RSA-AES-256-SHA • RSA-AES-256-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-AES-128-SHA • DHE-RSA-AES-256-SHA

PAN-OS 8.0 Cipher Suites Supported in FIPS-CC Mode

The following table lists cipher suites that are supported on firewalls running a PAN-OS® 8.0 release in FIPS-CC mode. The TOE includes NIST-validated cryptographic algorithms that provide support for the cryptographic functions listed in this table. The following functions have been certified in accordance with the identified standards. For certificate details, visit the Computer Security Resource Center (CSRC) on the National Institute of Standards and Technology (NIST) website (certificates in the table are linked to the corresponding page on the NIST website; once there, search for the certificate number to locate the detailed description).



If your firewall is running in normal (non-FIPS-CC) operational mode, see Cipher Suites Supported in PAN-OS 8.0

Functions	Standards	Certificates
Asymmetric key generation		
FFC key pair generation (key size 2048 bits)	FIPS PUB 186-4	**Appliances:** DSA #1207 **VMs:** DSA #1205
ECC key pair generation (NIST curves P-256, P-384)	FIPS PUB 186-4	**Appliances:** ECDSA #1103 **VMs:** ECDSA #1101
RSA key generation (key size 2048 bits)	FIPS PUB 186-4	**Appliances:** RSA #2467 **VMs:**

© 2018 Palo Alto Networks, Inc.

Functions	Standards	Certificates
		RSA #2463
Cryptographic Key Generation (for IKE Peer Authentication)		
RSA key generation (key size 2048 bits)	FIPS PUB 186-4	**Appliances:** RSA #2467 **VMs:** RSA #2463
ECDSAkey pair generation (NIST curves P-256, P-384)	FIPS PUB 186-4	**Appliances:** ECDSA #1103 **VMs:** ECDSA #1101
Cryptographic Key Establishment		
ECDSA based key establishment	NIST SP 800-56A Revision 2	**Appliances:** CVL #1211; CVL #1214 **VMs:** CVL #1206; CVL #1203
FFC based key establishment	NIST SP 800-56A Revision 2	**Appliances:** CVL #1211 **VMs:** CVL #1203
AES Data Encryption/Decryption		
• AES 128 and AES 256 • CBC 128 and CBC 256 • GCM 128 and GCM 256	• AES as specified in ISO 18033-3 • CBC as specified in ISO 10116 • GCM as specified in ISO 19772	**Appliances:** AES #4532 **VMs:** AES #4526
Signature Generation and Verification		
RSA Digital Signature Algorithm (rDSA) (modulus 2048)	FIPS PUB 186-4, "Digital Signature Standard (DSS)", Section 5.5, using PKCS #1 v2.1 Signature Schemes RSASSA-PSS and/or RSASSAPKCS1v1_5; ISO/IEC 9796-2, Digital signature scheme 2 or	**Appliances:** RSA #2467 **VMs:** RSA #2463

© 2018 Palo Alto Networks, Inc.

Functions	Standards	Certificates
	Digital Signature scheme 3	
ECDSA (NIST curves P-256, P-384, and P-521)	FIPS PUB 186-4, “Digital Signature Standard (DSS)”, Section 6 and Appendix D, Implementing "NIST curves" P-256, P-384, ISO/IEC 14888-3, Section 6.4	**Appliances:** RSA #2467 **VMs:** RSA #2463
Cryptographic hashing		
SHA-1, SHA-256, SHA-384 and SHA-512 (digest sizes 160, 256, 384 and 512 bits)	ISO/IEC 10118-3:2004	**Appliances:** SHS #3713 **VMs:** SHS #3707
Keyed-hash message authentication		
• HMAC-SHA-1 • MAC-SHA-256 • HMAC-SHA-384 • HMAC-SHA-512	ISO/IEC 9797-2:2011	**Appliances:** HMAC #2990 **VMs:** HMAC #2986
Random bit generation		
CTR_DRBG (AES) from a hardware based noise source with one independent software-based noise source of 256 bits of non-determinism	ISO/IEC 18031:2011	**Appliances:** DRBG #1489 **VMs:** DRBG #1486

© 2018 Palo Alto Networks, Inc.

Cipher Suites Supported in PAN-OS 7.1

The following topics list cipher suites that are supported on firewalls running a PAN-OS® 7.1 release in normal (non-FIPS-CC) operational mode.

The ciphers supported in normal operation mode are grouped according to feature or functionality in the following sections:

- PAN-OS 7.1 GlobalProtect Cipher Suites
- PAN-OS 7.1 IPSec Cipher Suites
- PAN-OS 7.1 IKE and Web Certificate Cipher Suites
- PAN-OS 7.1 Decryption Cipher Suites
- PAN-OS 7.1 Administrative Session Cipher Suites
- PAN-OS 7.1 PAN-OS-to-Panorama Connection Cipher Suites

PAN-OS 7.1 GlobalProtect Cipher Suites

The following table lists cipher suites for GlobalProtect™ that are supported on firewalls running a PAN-OS® 7.1 release in normal (non-FIPS-CC) operational mode.

- GlobalProtect App/Agent—SSL tunnels and SSL connections to gateway and portal
- GlobalProtect App/Agent—IPSec mode
- GlobalProtect Portal—Browser Access

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases
GlobalProtect App/Agent—SSL tunnels and SSL connections to gateway and portal	• TLSv1.0, TLSv1.1, and TLSv1.2 cipher suites • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-3DES-SHA • RSA-AES-128-SHA • RSA-AES-256-SHA • DHE-RSA-AES-128-SHA • DHE-RSA-AES-256-SHA • EDH-RSA-3DES-SHA • ECDHE-ECDSA-AES-128-SHA • ECDHE-ECDSA-AES-256-SHA • ECDHE-ECDSA-AES-128-GCM-SHA-256 • ECDHE-ECDSA-AES-256-GCM-SHA-384
GlobalProtect App/Agent—IPSec mode (Keys transported through SSL session with gateway)	• AES-128-CBC-HMAC-SHA • AES-128-GCM-HMAC-SHA • AES-256-GCM-HMAC-SHA
GlobalProtect Portal—Browser Access	• SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 cipher suites • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-3DES-SHA

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases
	• RSA-AES-128-SHA • RSA-AES-256-SHA • RSA-AES-128-SHA-256 • RSA-AES-256-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-AES-256-SHA • DHE-RSA-AES-128-SHA • EDH-RSA-3DES-SHA • ECDHE-ECDSA-AES-128-SHA • ECDHE-ECDSA-AES-256-SHA • ECDHE-ECDSA-AES-128-GCM-SHA-256 • ECDHE-ECDSA-AES-256-GCM-SHA-384

PAN-OS 7.1 IPSec Cipher Suites

The following table lists cipher suites for IPSec that are supported on firewalls running a PAN-OS® 7.1 release in normal (non-FIPS-CC) operational mode.

- IPSec—Encryption
- IPSec—Message Authentication
- IPSec—Key Exchange

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases
IPSec—Encryption	• 3DES • AES-128-CBC • AES-192-CBC • AES-256-CBC • AES-128-CCM • AES-128-GCM • AES-256-GCM
IPSec—Message Authentication	• NONE • HMAC-MD5 • HMAC-SHA • HMAC-SHA-256 • HMAC-SHA-384 • HMAC-SHA-512
IPSec—Key Exchange	Diffie-Hellman groups with or without perfect forward secrecy (PFS): • No PFS—This option specifies that the firewall reuses the same key for IKE phase 1 and phase 2 instead of renewing the key for phase 2. • Group 1 (768-bit keys) with PFS enabled • Group 2 (1024-bit keys) with PFS enabled • Group 5 (1536-bit keys) with PFS enabled • Group 14 (2048-bit keys) with PFS enabled

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases
	• Group 19 (256-bit elliptic curve group) with PFS enabled • Group 20 (384-bit elliptic curve group) with PFS enabled

PAN-OS 7.1 IKE and Web Certificate Cipher Suites

The following table lists cipher suites for Internet Key Exchange (IKE) and PAN-OS® web certificates that are supported on firewalls running a PAN-OS 7.1 release in normal (non-FIPS-CC) operational mode.

- IKE Certificate Support
- IKE—Encryption
- IKE—Message Authentication
- IKE—Key Exchange
- PAN-OS Web Certificates

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases

IKE Certificate Support	• RSA
	  • Keys—512-, 1024-, 2048-, and 3072-bit keys
	  • Digital signature algorithms—SHA, SHA-256, SHA-384, or SHA-512
	• ECDSA
	  • Keys—256- and 384-bit keys
	  • Digital signature algorithms—SHA-256, SHA-384, or SHA-512

IKE—Encryption	• 3DES • AES-128-CBC • AES-192-CBC • AES-256-CBC
IKE—Message Authentication	• HMAC-MD5 • HMAC-SHA • HMAC-SHA-256 • HMAC-SHA-384 • HMAC-SHA-512
IKE—Key Exchange	Diffie-Hellman groups • Group 1 (768-bit keys) • Group 2 (1024-bit keys) • Group 5 (1536-bit keys) • Group 14 (2048-bit keys) • Group 19 (256-bit elliptic curve group) • Group 20 (384-bit elliptic curve group)

PAN-OS Web Certificates	• RSA
	  • Keys—512-, 1024-, 2048-, 3072- and 4096-bit keys
	  • Digital signature algorithms—SHA, SHA-256, SHA-384, or SHA-512

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases

	• ECDSA
	  • Keys—256- and 384-bit keys
	  • Digital signature algorithms—SHA-256, SHA-384, or SHA-512

PAN-OS 7.1 Decryption Cipher Suites

The following table lists cipher suites for decryption that are supported on firewalls running a PAN-OS® 7.1 release in normal (non-FIPS-CC) operational mode.

- SSH Decryption (SSHv2 only)—Encryption
- SSH Decryption (SSHv2 only)—Message Authentication
- SSL/TLS Decryption
- SSL/TLS Elliptical Curves
- SSL/TLS Decryption—Perfect Forward Secrecy (PFS) Ciphers

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases
SSH Decryption (SSHv2 only)—Encryption	• AES-128-CBC • AES-192-CBC • AES-256-CBC • AES-128-CTR • AES-192-CTR • AES-256-CTR
SSH Decryption (SSHv2 only)—Message Authentication	• HMAC-RIPEMD • HMAC-MD5-96 • HMAC-MD5 • HMAC-SHA-96 • HMAC-RIPEMD-160 • HMAC-SHA
SSL/TLS Decryption *Firewalls do not decrypt sessions where Diffie-Hellman key exchange is used for Forward Proxy Decryption in PAN-OS 7.1 and earlier releases.*	• SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 cipher suites • RSA keys—512-, 1024-, 2048-, 3072-, 4096-, and 8192-bit keys *The firewall can authenticate certificates up to 8192-bit RSA keys from the destination server, however the firewall generated certificate to the client supports only up to 2048-bit RSA keys.* • RSA-RC4-128-MD5 • RSA-RC4-128-SHA • RSA-3DES-EDE-CBC-SHA • RSA-AES-128-CBC-SHA • RSA-AES-256-CBC-SHA • RSA-AES-128-CBC-SHA-256 • RSA-AES-256-CBC-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases
SSL/TLS Elliptical Curves	PAN-OS 7.1 does not support elliptical curve cryptography (ECC) certificates for SSL/TLS decryption.
SSL/TLS Decryption—Perfect Forward Secrecy (PFS) Ciphers (SSL Forward Proxy Decryption only)	• DHE-RSA-3DES-EDE-CBC-SHA • DHE-RSA-AES-128-CBC-SHA • DHE-RSA-AES-256-CBC-SHA • DHE-RSA-AES-128-GCM-SHA-256 • DHE-RSA-AES-256-GCM-SHA-384 • ECDHE-RSA-AES-128-CBC-SHA • ECDHE-RSA-AES-256-CBC-SHA • ECDHE-RSA-AES-128-GCM-SHA-256 • ECDHE-RSA-AES-256-GCM-SHA-384 Additional ciphers supported in PAN-OS 7.1.13 and later PAN-OS 7.1 releases: • DHE-RSA-AES-128-CBC-SHA-256 • DHE-RSA-AES-256-CBC-SHA-256 • ECDHE-RSA-AES-128-CBC-SHA-256 • ECDHE-RSA-AES-256-CBC-SHA-384

PAN-OS 7.1 Administrative Session Cipher Suites

The following table lists the cipher suites for administrative sessions that are supported on firewalls running a PAN-OS® 7.1 release in normal (non-FIPS-CC) operational mode.

- Administrative Sessions to Web Interface
- Administrative Sessions to CLI (SSH)—Encryption
- Administrative Sessions to CLI (SSH)—Message Authentication
- Administrative Sessions to CLI (SSH)—Server Host Key Types
- Administrative Sessions to CLI (SSH)—Key Exchange Algorithms

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases
Administrative Sessions to Web Interface	• SSLv3, TLSv1.0, TLSv1.1 and TLSv1.2 cipher suites • RSA-3DES-EDE-CBC-SHA • RSA-SEED-CBC-SHA • RSA-CAMELLIA-128-CBC-SHA • RSA-CAMELLIA-256-CBC-SHA • RSA-AES-128-SHA • RSA-AES-128-SHA256 • RSA-AES-256-SHA • RSA-AES-256-SHA256 • RSA-AES-128-GCM-SHA256 • RSA-AES-256-GCM-SHA384 • DHE-RSA-3DES-EDE-SHA • DHE-RSA-AES-128-SHA • DHE-RSA-AES-256-SHA

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases
	• ECDHE-ECDSA-AES-128-SHA • ECDHE-ECDSA-AES-256-SHA • ECDHE-ECDSA-AES-128-GCM-SHA256 • ECDHE-ECDSA-AES-256-GCM-SHA384
Administrative Sessions to CLI (SSH)—Encryption	• 3DES • ARCFOUR128 • ARCFOUR256 • BLOWFISH • CAST128 • AES-128-CBC • AES-192-CBC • AES-256-CBC • AES-128-CTR • AES-192-CTR • AES-256-CTR • AES-128-GCM • AES-256-GCM
Administrative Sessions to CLI (SSH)—Message Authentication	• UMAC-64 • HMAC-MD5-96 • HMAC-MD5 • HMAC-SHA-96 • HMAC-RIPEMD-160 • HMAC-SHA • HMAC-SHA-256 • HMAC-SHA-512
Administrative Sessions to CLI (SSH)—Server Host Key Types	• RSA keys—2048-bit keys • DSA keys—1024-bit keys
Administrative Sessions to CLI (SSH)—Key Exchange Algorithms	• diffie-hellman-group1-sha1 • diffie-hellman-group14-sha1 • diffie-hellman-group-exchange-sha1 • diffie-hellman-group-exchange-sha256 • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521

PAN-OS 7.1 PAN-OS-to-Panorama Connection Cipher Suites

The following table lists the cipher suites for PAN-OS®-to-Panorama™ connections that are supported on firewalls running a PAN-OS 7.1 release in normal (non-FIPS-CC) operational mode.

© 2018 Palo Alto Networks, Inc.

Feature or Function	Ciphers Supported in PAN-OS 7.1 Releases
PAN-OS to Panorama Connection	• RSA-RC4-128-SHA • RSA-3DES-SHA • RSA-SEED-SHA • RSA-CAMELLIA-128-SHA • RSA-CAMELLIA-256-SHA • RSA-AES-128-SHA • RSA-AES-128-SHA-256 • RSA-AES-256-SHA • RSA-AES-256-SHA-256 • RSA-AES-128-GCM-SHA-256 • RSA-AES-256-GCM-SHA-384 • DHE-RSA-AES-128-SHA • DHE-RSA-AES-256-SHA

© 2018 Palo Alto Networks, Inc.

GlobalProtect

The following topics provide support information for the GlobalProtect™ app (originally referred to as the GlobalProtect agent on Windows and Mac).

- Where Can I Install the GlobalProtect App?
- What X-Auth IPSec Clients are Supported?
- What Features Does GlobalProtect Support?
- What GlobalProtect Features Do Third-Party Mobile Device Management Systems Support? on page 83

© 2018 Palo Alto Networks, Inc.

Where Can I Install the GlobalProtect App?

The following table shows operating systems on which you can install each release of the GlobalProtect app.

Operating System	Release 4.0	Release 4.1	Release 5.0
Apple iOS			
iOS 8	–	–	–
iOS 9	–	–	–
iOS 10	✓ (64-bit devices only)	✓ (64-bit devices only)	✓ (64-bit devices only)
iOS 11	–	✓ (64-bit devices only)	✓ (64-bit devices only)
iOS 12	–	–	✓ (64-bit devices only)
Apple Mac			
Mac OS X 10.5 (64-bit only)	–	–	–
Mac OS X 10.6	–	–	–
Mac OS X 10.7	–	–	–
Mac OS X 10.8	–	–	–
Mac OS X 10.9	✓	–	–
Mac OS X 10.10	✓	✓	–
macOS 10.11	✓	✓	–
macOS 10.12	✓	✓	–
macOS 10.13	✓ *4.0.3 & later	✓	–

© 2018 Palo Alto Networks, Inc.

Operating System	Release 4.0	Release 4.1	Release 5.0
macOS 10.14	–	✓ *4.1.5 & later	–
Google Android			
Google Android 4.4	✓	✓	–
Google Android 5.x	✓	✓	–
Google Android 6.x	✓	✓	–
Google Android 7.x	✓	✓	–
Google Android 8.x	–	✓	–
Google Android 9.x	–	✓ *4.1.5 & later	–
Google Chrome			
Google Chrome OS 45 and later releases	✓	✓	–
Linux			
CentOS 7	–	✓	–
Red Hat Enterprise Linux (RHEL) 7	–	✓	–
Ubuntu 14.04 and later releases	–	✓	–
Ubuntu 16	–	✓	–
Windows			
Windows XP (32-bit only)	–	–	–
Windows Vista	–	–	–
Windows 7	✓	✓	–
Windows 8	✓	✓	–
Windows 8.1	✓	✓	–

© 2018 Palo Alto Networks, Inc.

Operating System	Release 4.0	Release 4.1	Release 5.0
Windows 10	✓	✓	—
Windows 10 UWP	✓	✓	—

© 2018 Palo Alto Networks, Inc.

What X-Auth IPSec Clients are Supported?

The following table lists third-party X-Auth IPSec client support for each PAN-OS® release.

Third-Party X-Auth IPSec Client	Minimum PAN-OS Release
iOS built-in IPSec client	6.1
Android built-in IPSec client	6.1
VPNC on Ubuntu Linux 10.04 and later versions and CentOS 6 and later versions	6.1
strongSwan on Ubuntu Linux and CentOS*	6.1
* To set up authentication for strongSwan Ubuntu and CentOS clients for PAN-OS 7.1 and later releases, refer to the GlobalProtect Administrator's Guide for yourrelease.	

© 2018 Palo Alto Networks, Inc.

What Features Does GlobalProtect Support?

The following table lists the features supported on GlobalProtect by OS. An entry in the table indicates the first supported release of the feature on the OS. A – indicates the feature is not supported. For recommended minimum GlobalProtect app versions, see *What OS Versions are Supported with GlobalProtect?* in the GlobalProtect Administrator's Guide for your PAN-OS version.

Feature	Android	iOS	Chrome	Windows	Windows 10 UWP	Mac	Linux
Authentication							
App login enhancements	4.0.0	4.0.0	–	4.0.0	–	4.0.0	–
Multi-factor authentication policy	–	–	–	4.0.0	–	4.0.0	–
SAML authentication	4.0.0	4.0.0 (**On-Demand** connect method only)	4.1.0	4.0.0	–	4.0.0	–
Expired Active Directory password change for remote users	4.1.0	4.1.0 (notifications only) 5.0.0 (full support)	4.1.0	4.1.0	4.1.0	4.1.0	–
Active Directory password change using the GlobalProtect Credential Provider	–	–	–	4.1.0	–	–	–
Single Sign-On (SSO)							
SSO (Credential Provider)	–	–	–	1.2.0	–	–	–
Kerberos SSO	–	–	–	3.0.0	–	4.1.0	–
VPN Connections							

© 2018 Palo Alto Networks, Inc.

Feature	Android	iOS	Chrome	Windows	Windows 10 UWP	Mac	Linux
IPSec	1.3.0	1.3.0	3.1.1	1.0.0	4.0.5	1.0.0	4.1
SSL	1.3.0	1.3.0	3.1.1	1.0.0	3.1.3	1.0.0	4.1
Clientless VPN	– (no client required)	– (no client required)	– (no client required)	– (no client required)	– (no client required)	– (no client required)	– (no client required)
Connect Methods							
User-logon (always on)	1.3.0	1.3.0	–	1.0.0	3.1.3 (Always On configured from third-party MDM)	1.0.0	4.1
Pre-logon (always-on)	–	–	–	1.1.0	–	1.1.0	–
Pre-logon (then on-demand)	–	–	–	3.1.0	–	3.1.0	–
On-demand	1.3.0	1.3.0	3.1.1	1.0.0	3.1.3	1.0.0	4.1
Connection Priority							
External gateway priority by source region	4.0.0	4.0.0	4.0.0	4.0.0	4.0.0	4.0.0	4.1
Internal gateway selection by source IP address	4.0.0 (Except DHCP options)	4.0.0 (Except DHCP options)	–	4.0.0	–	4.0.0	4.1
Modes							
Internal mode	1.3.0	1.3.0	–	1.0.0	–	1.0.0	4.1
External mode	1.3.0	1.3.0	3.1.1	1.0.0	3.1.3	1.0.0	4.1
Networking							
IPv4 addressing	1.3.0	1.3.0	3.1.1	1.0.0	3.1.3	1.0.0	4.1

© 2018 Palo Alto Networks, Inc.

Feature	Android	iOS	Chrome	Windows	Windows 10 UWP	Mac	Linux
IPv6 addressing	4.0.0	4.0.0	4.0.0	4.0.0	4.0.0	4.0.0	4.1
Split tunnel based on access route	–	4.0.0	4.0.0	4.0.0	4.0.0	4.0.0	4.1
Split tunnel based on destination domain, client process, and video streaming application	–	–	–	4.1.0	–	4.1.0	–
Customization							
GlobalProtect Credential Provider pre-logon connection status	–	–	–	4.1.0	–	–	–
Multiple portal support	–	–	–	4.1.0	–	4.1.0	–
Resilient VPN	4.0.3	4.0.3	–	4.0.3	–	4.0.3	–
Pre-logon tunnel rename timeout	–	–	–	4.0.2	–	–	–
Restrict transparent app upgrades to internal network connections	–	–	–	4.0.0	–	4.0.0	–
Enforce GlobalProtect for network access	–	–	–	3.1.0	3.1.3 (VPN Lockdown configured from third-party MDM)	3.1.0	–
Deployment of SSL Forward	–	–	–	3.0.0	–	3.0.0	–

© 2018 Palo Alto Networks, Inc.

Feature	Android	iOS	Chrome	Windows	Windows 10 UWP	Mac	Linux
Proxy CA certificates in the trust store							
HIP reports	1.3.0	1.3.0	3.0.0	1.0.0	3.1.3 (Host information only; Notifications not supported)	1.0.0	4.1 (Host information only)
Run scripts before and after sessions	–	–	–	2.3.0	–	2.3.0	–
Certificate selection by OID	–	–	–	3.0.0	–	3.0.0	–
Allow users to disable GlobalProtect	–	–	–	2.2.0	–	2.2.0	4.1
Welcome and help pages	1.3.0	1.3.0	3.0.0	1.0.0	–	1.0.0	–
Other							
Automatic VPN Reconnect for Chromebooks	–	–	4.1.0	–	–	–	–

© 2018 Palo Alto Networks, Inc.

What GlobalProtect Features Do Third-Party Mobile Device Management Systems Support?

The following table lists the GlobalProtect features supported on third-party mobile device management (MDM) systems. A – indicates the feature is not supported.

Feature	AirWatch	Microsoft Intune	MobileIron	Google Admin Console
GlobalProtect app deployment	✓	✓	✓	✓
Always on VPN Configuration	✓	✓ (iOS and Windows 10 UWP only)	✓ (iOS and Android only)	–
Remote access VPN configuration	✓	✓ (iOS only)	✓ (iOS only)	–
Per-app VPN configuration	✓	✓ (iOS and Windows 10 UWP only)	✓ (iOS only)	–
MDM integration with HIP	✓	–	–	–
VPN lockdown	✓	–	–	–

© 2018 Palo Alto Networks, Inc.

User-ID Agent

You install the User-ID™ agent on a domain server that is running a supported operating system (OS) and then connect the User-ID agent to exchange or directory servers.

> Where Can I Install the User-ID Agent? on page 87
> Which Servers Can the User-ID Agent Monitor? on page 88

© 2018 Palo Alto Networks, Inc.

Where Can I Install the User-ID Agent?

The following table shows the operating systems on which you can install each release of the Windows-based User-ID agent.

Operating System	Release 6.0	Release 7.0	Release 8.0	Release 8.1
Windows Server 2003	✓	✓	–	–
Windows Server 2003 R2	✓	✓	–	–
Windows Server 2008 and 2008 R2	✓	✓	✓	✓
Windows Server 2012	✓	✓	✓	✓
Windows Server 2012 R2	✓ *6.0.3 & later	✓	✓	✓
Windows Server 2016	–	–	✓ *8.0.5 & later	✓

© 2018 Palo Alto Networks, Inc.

Which Servers Can the User-ID Agent Monitor?

The following are the exchange and directory servers you can monitor with the PAN-OS integrated and Windows-based User-ID agents:

Server	Versions Supported

Microsoft Exchange Server	• 2016—with Windows User-ID agent 8.0.5 and later releases or with PAN-OS integrated User-ID agents running the following PAN-OS releases:
	  • PAN-OS 8.0.4 and later releases
	  • PAN-OS 7.1.12 and later PAN-OS 7.1 releases
	  • PAN-OS 7.0.18 and later PAN-OS 7.0 releases
	  • PAN-OS 6.1.18 and later PAN-OS 6.1 releases
	• 2013
	• 2010 (14.0)
	• 2007 (8.0)

Microsoft Windows Server	*You can install only specific releases of the Windows-based User-ID agent on supported Windows server versions.*
	• 2016—only with Windows User-ID agent 8.0.5 and later releases or with PAN-OS integrated User-ID agents running the following PAN-OS releases:
	  • PAN-OS 8.0.4 and later releases
	  • PAN-OS 7.1.12 and later PAN-OS 7.1 releases
	  • PAN-OS 7.0.18 and later PAN-OS 7.0 releases
	  • PAN-OS 6.1.18 and later PAN-OS 6.1 releases
	• 2012 and 2012 R2
	• 2008 and 2008 R2
	• 2003 and 2003 R2

Novell eDirectory Server	8.8

Terminal Services (TS) Agent

You install the Terminal Services (TS) agent on a domain server that is running a supported operating system (OS) and then report username-to-port mapping information to PAN-OS firewalls.

> Where Can I Install the Terminal Services (TS) Agent?
> How Many TS Agents Does My Firewall Model Support? on page 92

© 2018 Palo Alto Networks, Inc.

Where Can I Install the Terminal Services (TS) Agent?

The following table shows the operating systems on which you can install each release of the TS agent.

Operating System	TS agent 6.0 Release	TS agent 7.0 Release	TS agent 8.0 Release	TS agent 8.1 Release
Windows Server 2003	✓	–	–	–
Windows Server 2003 R2	✓	–	–	–
Windows Server 2008	✓	✓	✓	✓
Windows Server 2008 R2	✓	✓	✓	✓
Windows Server 2012 R2	✓ *6.0.3 & later	✓	✓	✓
Windows Server 2016	–	–	–	✓ *8.1.1 & later
Citrix Metaframe Presentation Server 4.x	✓	✓	✓	✓
Citrix XenApp 5.x	✓	✓	✓	✓
Citrix XenApp 6.x	✓	✓	✓	✓
Citrix XenApp 7.x	✓ *6.0.4 & later	✓	✓	✓

© 2018 Palo Alto Networks, Inc.

How Many TS Agents Does My Firewall Model Support?

The following table shows how many Terminal Services (TS) agents each hardware-based and VM-Series firewall supports.



For optimal configuration, install the TS agent version that matches the PAN-OS version running on the firewall. If there is not a TS agent version that matches the PAN-OS version, install the latest version that is closest to the PAN-OS version.

Firewall Model or VM	PAN-OS 7.1, 8.0, and 8.1 Release
PA-7080, PA-7050	1000
PA-5200 Series	1000
PA-3200 Series	400
PA-3000 Series	400
PA-850	400
PA-220	400
PA-220R	400
VM-700	1000
VM-300	400
VM-100	400
VM-50 Lite	400

Traps

You can install the Traps™ agent and the Endpoint Security Manager (ESM) Components (comprised of the ESM Console, one or more ESM Servers, and the database) only on servers and endpoints that are running a supported operating system (OS).

> Where Can I Install the Endpoint Security Manager (ESM)?
> Where Can I Install the Traps Agent?

© 2018 Palo Alto Networks, Inc.

Where Can I Install the Endpoint Security Manager (ESM)?

The Endpoint Security Manager (ESM) comprises the ESM Console, one or more ESM Servers, and a database. You can install the ESM components on dedicated servers or install them on the same server as long as you install them on a supported operating system (OS).

Server Operating System	ESM 3.4	ESM 4.1	ESM 4.2
Windows Server 2008 R2	✓	✓	✓
Windows Server 2012	✓	✓	✓
Windows Server 2012 R2	✓	✓	✓
Windows Server 2016	–	✓	✓

© 2018 Palo Alto Networks, Inc.

Where Can I Install the Traps Agent?

Palo Alto Networks strives to support the latest major operating systems. After a new operating system is available, there may be a short delay in our support of the operating system as we test interoperability with the production release.

You can install the Traps agent on supported hardware and virtual endpoints:

- Endpoint Operating Systems Supported with Traps
- Virtual Desktop Operating Systems Supported with Traps
- Virtual Applications Supported with Traps

Endpoint Operating Systems Supported with Traps

The following table shows the endpoint operating systems on which you can install each release of the Traps agent. These operating systems are also supported with supported Citrix and VMware virtual applications (see Virtual Applications Supported with Traps).

Endpoint Operating System	Traps 3.4	Traps 4.1	Traps 4.2	Traps 5.0
Android				
4.4	–	–	–	✓
5	–	–	–	✓
6	–	–	–	✓
7	–	–	–	✓
8	–	–	–	✓
Linux				
Amazon Linux 2 LTS Candidate (2017.12)	–	–	✓ *4.2.1 & later	✓ *5.0.1 & later
Amazon Linux 2 LTS Candidate 2	–	–	✓ *4.2.1 & later	✓ *5.0.1 & later
CentOS 6	–	–	✓	✓
CentOS 7	–	–	✓	✓
Debian 8	–	–	✓	✓

© 2018 Palo Alto Networks, Inc.

Endpoint Operating System	Traps 3.4	Traps 4.1	Traps 4.2	Traps 5.0
				*5.0.1 & later
Debian 9	–	–	✓	✓ *5.0.1 & later
Oracle 6	–	–	✓	✓ *5.0.1 & later
Oracle 7	–	–	✓	✓ *5.0.1 & later
Red Hat Enterprise Linux 6	–	–	✓	✓
Red Hat Enterprise Linux 7	–	–	✓	✓
SUSE Linux Enterprise Server 12	–	–	✓	✓
Ubuntu Server 12	–	–	✓	✓
Ubuntu Server 14	–	–	✓	✓
Ubuntu Server 16	–	–	✓	✓
Mac				
OS X 10.10	–	–	–	–
OS X 10.11	–	✓	✓	✓
macOS 10.12	–	✓	✓	✓
macOS 10.13	–	✓ *4.1.1 & later	✓	✓
macOS 10.14	–	✓ *4.1.5-h1 & later	✓ *4.2.1-h3 & later	✓ *5.0.3-h1 & later
Windows				
Windows 2000 and below	–	–	–	–

© 2018 Palo Alto Networks, Inc.

Endpoint Operating System	Traps 3.4	Traps 4.1	Traps 4.2	Traps 5.0
Windows XP (32-bit) (SP2 and earlier)	–	–	–	–
Windows XP (32-bit) * (SP3 and later)	✓	✓	✓	✓
Windows XP (64 bit)	–	–	–	–
Windows Vista (SP1 and later; FIPS mode)	✓	✓	✓	✓
Windows 7 (RTM and SP1; FIPS mode; all editions except Home) SHA-2 code signing supported with Microsoft Security Advisory 3033929	✓	✓	✓	✓
Windows Embedded 7 (Standard and POSReady)	✓	✓	✓	✓
Windows 8.1 (FIPS mode)	✓	✓	✓	✓
Windows Embedded 8.1 Pro	✓	✓	✓	✓
Windows Embedded POSReady 2009	–	✓ *4.1.2 & later	✓	✓
Windows 10 Education	–	–	–	✓ *5.0.4 & later
Windows 10 Pro (CB and CBB)	✓	✓	✓	✓
Windows 10 Enterprise (CB, CBB, and LTSB)	✓	✓	✓	✓

© 2018 Palo Alto Networks, Inc.

Endpoint Operating System	Traps 3.4	Traps 4.1	Traps 4.2	Traps 5.0
Do not upgrade Windows 10 to build 1703 when running Traps 4.0.0; instead, you must first upgrade to Traps 4.0.1 or a later release before you upgrade to Windows 10 1703.				
Windows 10 Fall Creators Update 1709	–	✓ *4.1.2 & later	✓	✓
Windows 10 Update 1803	–	✓ *4.1.4 & later	✓	✓
Windows 10 Update 1809	–	–	–	✓ *5.0.4 & later
Windows Server 2003 (SP1 and earlier)	–	–	–	–
Windows Server 2003 (32-bit) (SP2 and later)	✓	✓	✓	✓
Windows Server 2003 (64-bit)	–	–	–	–

© 2018 Palo Alto Networks, Inc.

Endpoint Operating System	Traps 3.4	Traps 4.1	Traps 4.2	Traps 5.0
Windows Server 2003 R2 (32-bit) (SP2 and later)	✓	✓	✓	✓
Windows Server 2003 R2 (64 bit)	–	–	–	–
Windows Server Core option (Windows Server 2012, 2012 R2, and 2016 only)	✓ *3.4.2 & later	✓	✓	✓
Windows Server 2008 (FIPS mode)	✓	✓	✓	✓
Windows Server 2008 R2 (FIPS mode)	✓	✓	✓	✓
Windows Server 2012 (All editions; FIPS mode)	✓	✓	✓	✓
Windows Server 2012 R2 (All editions; FIPS mode)	✓	✓	✓	✓
Windows Server 2016 (Standard edition; Server with Desktop Experience - previously known as Server with a GUI)	✓ *3.4.2 & later	✓	✓	✓
Windows Server 2016 Datacenter edition	–	–	–	✓ *5.0.4 & later

* Palo Alto Networks supports Microsoft operating systems for five years after the official Microsoft end-of-life (EoL) date with one exception: we will continue support for Windows XP—only as indicated in this table—until support ends for Windows 2003 (currently July 2020).

© 2018 Palo Alto Networks, Inc.

Virtual Desktop Operating Systems Supported with Traps

The following table shows the minimum supported versions of virtual desktop operating systems on which you can install each release of the Traps agent.

Product	Traps 3.4	Traps 4.1	Traps 4.2	Traps 5.0
Citrix				
Citrix Provisioning Services (PVS)	7.6	7.13	7.15	7.15
Citrix XenDesktop RDS	7.6	7.13	7.15	7.15
Citrix XenDesktop VDI	7.6	7.13	7.15	7.15
VMware				
VMware Horizon View RDS	6.2	7.1	7.2	7.2
VMware Horizon View VDI	6.2	7.1	7.2	7.2
Windows				
Windows Server R2 RDS	2008	2008	2008	2008
Windows Server 2012 R2 RDS	2012	2012	2012	2012
Windows Server 2016 RDS	—	2016	2016	2016

Virtual Applications Supported with Traps

The following table shows the supported software vendors you can use to deploy virtual applications and the minimum software version supported with each release of the Traps agent.

Product	Traps 3.4	Traps 4.1	Traps 4.2	Traps 5.0
Citrix				
Citrix XenApp	7.6	7.13	7.6 and 7.15	7.6 and 7.15
Citrix App Layering (supports Traps installed on the OS layer)	4	4	4	4

© 2018 Palo Alto Networks, Inc.

Product	Traps 3.4	Traps 4.1	Traps 4.2	Traps 5.0
VMware				
VMware AppVolumes	2.6	2.12 *Traps 4.1.2 and later	2.13.1	2.13.1
VMware ThinApp	5.2	5.2.2	5.2.2	5.2.2
Windows				
Windows Virtual PC	1.3.7600.16432 *Traps 3.4.4 only	—	—	—

* Windows 8.1 and later releases with AppVolumes does not support the anti-ransomware protection module. See the Traps 4.1 Release Notes for details.

IPv6 Support by Feature

© 2018 Palo Alto Networks, Inc.

IPv6 Support by Feature

Use the following table to review PAN-OS® features (listed by category) that support IPv6 traffic.

- Security
- Management & Panorama
- Networking
- VPN
- Host Dynamic Address Configuration
- Device
- User-ID

PAN-OS Feature	PAN-OS 6.1	PAN-OS 7.1	PAN-OS 8.0	PAN-OS 8.1
Security				
App-ID™ and Firewalling in Layer 2 and Layer 3	✓	✓	✓	✓
User-ID™	✓	✓	✓	✓
Content-ID™	✓	✓	✓	✓
Block IPv6 in IPv4 tunneling (via App-ID)	✓	✓	✓	✓
URL Filtering	✓	✓	✓	✓
SSL Decryption	✓	✓	✓	✓
SSH Decryption	✓	✓	✓	✓
DoS Rulebase	✓	✓	✓	✓
Management & Panorama				
SSH Management (dedicated MGMT port)	✓	✓	✓	✓
Web Interface Management (dedicated MGMT port)	✓	✓	✓	✓
Interface Management (ping, telnet, ssh, http, https - all ports)	✓	✓	✓	✓
Device to Panorama SSL TCP connection	✓	✓	✓	✓
Panorama HA connection between peers	–	–	✓	✓

© 2018 Palo Alto Networks, Inc.

PAN-OS Feature	PAN-OS 6.1	PAN-OS 7.1	PAN-OS 8.0	PAN-OS 8.1
DNS	✓	✓	✓	✓
RADIUS	✓	✓	✓	✓
LDAP	✓	✓	✓	✓
SYSLOG	✓	✓	✓	✓
SNMP	✓	✓	✓	✓
NTP	✓	✓	✓	✓
Device DNS (device only)	✓	✓	✓	✓
DNS proxy	✓	✓	✓	✓
Reporting and Visibility into IPv6	✓	✓	✓	✓
IPv6 address objects	✓	✓	✓	✓
IPv6 FQDN address objects	✓	✓	✓	✓
Networking				
IPv6 static routes	✓	✓	✓	✓
PBF	✓	✓	✓	✓
PBF next hop monitor (v6 endpoint)	✓	✓	✓	✓
OSPFv3	✓	✓	✓	✓
MP-BGP	–	–	✓	✓
ECMP	–	✓	✓	✓
Dual stack support for L3 interfaces	✓	✓	✓	✓
QoS Policy	✓	✓	✓	✓
QoS Marking	✓	✓	✓	✓
DSCP (session based)	–	✓	✓	✓

© 2018 Palo Alto Networks, Inc.

PAN-OS Feature	PAN-OS 6.1	PAN-OS 7.1	PAN-OS 8.0	PAN-OS 8.1
Neighbor Discovery and Duplicate Address Detection	✓	✓	✓	✓
Tunnel Content Inspection	–	–	✓	✓
Zone protection	✓	✓	✓	✓
NPTv6 (stateless prefix translation)	–	✓	✓	✓
NAT64 (IP-IPv6 protocol translation)	✓	✓	✓	✓
LLDP (Link Layer Discovery Protocol)	–	✓	✓	✓
Bidirectional Forwarding Detection (BFD)	–	✓	✓	✓
VPN				
GlobalProtect™	–	–	✓	✓
IKE/IPSec	–	✓	✓	✓
IKEv2	–	✓	✓	✓
IPv6 over IPv4 IPSec tunnel	✓	✓	✓	✓
Large Scale VPN (LSVPN)	–	–	✓	✓
Host Dynamic Address Configuration				
DHCPv6 relay	✓	✓	✓	✓
SLAAC (Router Advertisements)	✓	✓	✓	✓
SLAAC (Router Preference)	✓	✓	✓	✓
SLAAC (RDNSS)	–	–	✓	✓
Device				
High Availability (HA) - Active/Active	✓	✓	✓	✓
High Availability (HA) - Active/Passive	✓	✓	✓	✓
High Availability (HA) - IPv6 transport for HA1 & HA2	✓	✓	✓	✓

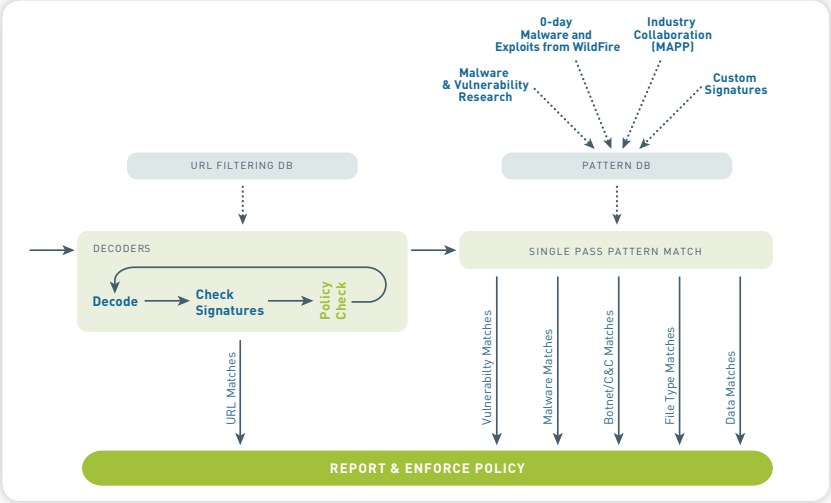
© 2018 Palo Alto Networks, Inc.

PAN-OS Feature	PAN-OS 6.1	PAN-OS 7.1	PAN-OS 8.0	PAN-OS 8.1
High Availability (HA) Path Monitoring (IPv6 Endpoint)	✓	✓	✓	✓
User-ID				
Map IPv6 Address to Users	✓	✓	✓	✓
Captive Portal for IPv6	✓	✓	✓	✓
Connection to User-ID Agents over IPv6	✓	✓	✓	✓
User-ID XML API for IPv6	✓	✓	✓	✓
Terminal Server Agent IPv6	✓	✓	✓	✓

Content-ID™

Content-ID enables customers to apply policies to inspect and control content traversing the network.

- Detect and block known and unknown threats in a single pass.
- Implement policy control over unapproved web surfing.
- Limit unauthorized transfer of files and sensitive data such as CC# or SSN.
- Proactively identify and defend against unknown, new or custom malware and exploits.
- Single pass software architecture maximizes performance by scanning traffic only once, regardless of which Content-ID features are enabled.



How Content-ID works.

Content-ID combines a real-time threat prevention engine with a comprehensive URL database and elements of application identification to limit unauthorized data and file transfers, detect and block a wide range of exploits, malware, dangerous web surfing as well as targeted and unknown threats. The application visibility and control delivered by App-ID™, combined with the content inspection enabled by Content-ID means that IT departments can regain control over application traffic and related content.

Enterprises of all sizes are at risk from a variety of increasingly sophisticated network-borne threats that have evolved to avoid many of the industry's traditional security measures. Palo Alto Networks® Content-ID delivers a new approach based on the complete analysis of all allowed traffic using multiple threat prevention and data-loss prevention techniques in a single unified engine. Unlike traditional solutions, Palo Alto Networks actually controls the threat vectors themselves through the tight control of all types of applications. This immediately reduces the "attack surface" of the network after which all allowed traffic is analyzed for exploits, malware, dangerous URLs, dangerous or restricted files or content. Palo Alto Networks then goes beyond stopping known threats to proactively identify and control unknown malware and exploits, which are often used as the leading edge of sophisticated network attacks.

Content-ID is built on a single-pass architecture, which is a unique integration of software and hardware that simplifies management, streamlines processing and maximizes performance. The single-pass architecture (SP3) integrates multiple threat prevention disciplines (IPS, anti-malware, URL filtering, etc.) into a single stream-based engine with a uniform signature format. This allows traffic to be fully analyzed in a single pass without the incremental performance degradation seen in other multi-function gateways. The software is tied directly to a parallel processing hardware platform that uses function specific processors for threat prevention to maximize throughput and minimize latency.

paloalto networks®

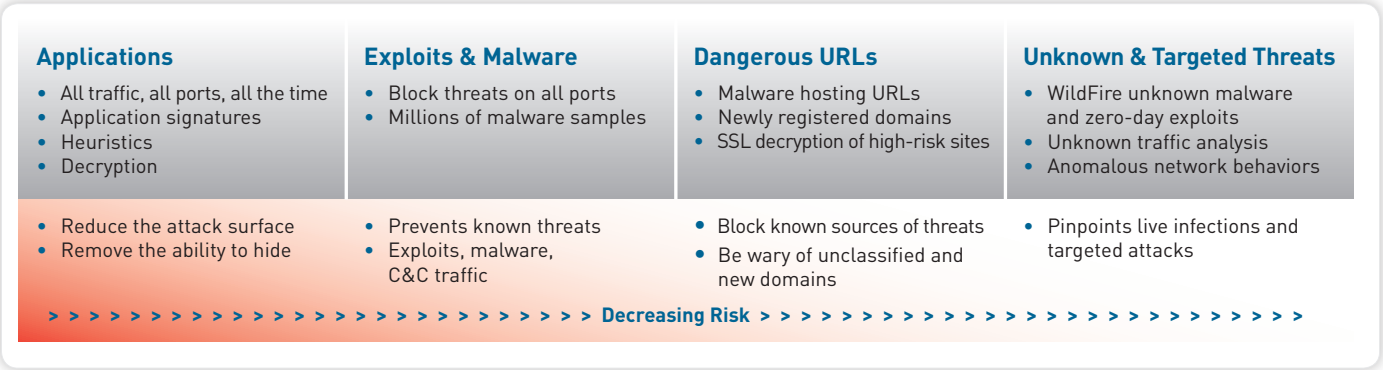
Threat Prevention

Enterprise networks are facing a rapidly evolving threat landscape full of modern applications, exploits, malware and attack strategies that are capable of avoiding traditional methods of detection. Threats are delivered via applications that dynamically hop ports, use non-standard ports, tunnel within other applications or hide within proxies, SSL or other types of encryption. These techniques can prevent traditional security solutions such as IPS and firewalls from ever inspecting the traffic, thus enabling threats to easily and repeatedly flow across the network. Additionally, enterprises are exposed to targeted and customized malware, which may pass undetected through traditional antivirus solutions.

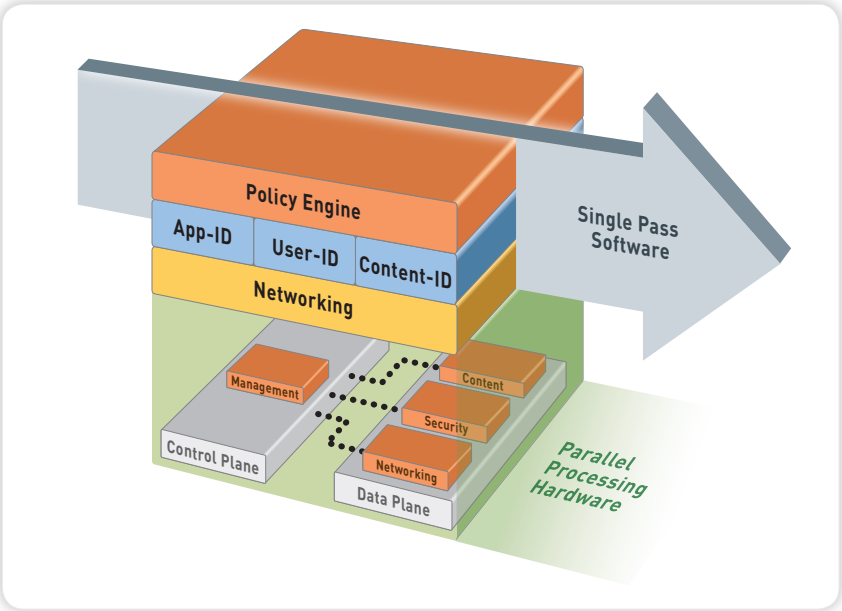
Palo Alto Networks Content-ID addresses these challenges with unique threat prevention abilities not found in other security solutions. First, the next-generation firewall removes the methods that threats use to hide from security through the complete analysis of all traffic, on all ports regardless of evasion, tunneling or circumvention techniques. Simply put, no threat prevention solution will be effective if it does not have visibility into the traffic, and only Palo Alto Networks ensures that visibility through the identification and control of all traffic.

- **Application decoders:** Content-ID leverages the more than 100 application and protocol decoders in App-ID to look for threats hidden within streams of application data. This enables the firewall to detect and prevent threats tunneled within approved applications that would pass by traditional IPS or proxy solutions.
- **SSL decryption:** More and more web traffic is encrypted with SSL by default. This can provide some protection to end-users, but it also can provide attackers with an encrypted channel to deliver exploits and malware. Palo Alto Networks ensures visibility by giving security organizations the flexibility to, by policy, granularly look inside of SSL traffic based on application or URL category.
- **Control of circumventing technologies:** Attackers and malware have increasingly turned to proxies, anonymizers and a variety of encrypted proxies to hide from traditional network security products. Palo Alto Networks provides the ability to tightly control these technologies and limit them to approved users, while blocking unapproved communications that could be used by attackers.
- **Uniform threat signature format:** Rather than use a separate set of scanning engines and signatures for each type of threat, Content-ID leverages a uniform threat engine and signature format to detect and block a wide range of malware including viruses, spyware, and vulnerability exploits in a single pass.

Secondly, Palo Alto Networks brings multiple security disciplines into a single context and a single threat prevention engine. This context enables security teams to easily see beyond individual security events and recognize the full extent of a threat. Security managers can now see the interconnection of applications, exploits, malware, URLs, DNS, anomalous network behaviors and management and reporting and ensures predictable performance by analyzing traffic once instead of progressive scanning in multiple engines.



Complete Threat Prevention
Palo Alto Networks brings a unique approach to Threat Prevention, starting with reducing the attack surface through positive security controls, blocking all known threats, and quickly discovering and stopping unknown threats.



Palo Alto Networks single pass parallel processing architecture accelerates content inspection performance while minimizing latency.

Integrated by Design

Palo Alto Networks next-generation firewalls are purpose-built platforms that utilize a single pass parallel processing architecture to maximize throughput and minimize latency. Traditional blade or UTM architectures notoriously introduce performance penalties for each feature that is enabled due to repeatedly processing traffic for each blade or feature. Palo Alto Networks designed a unique approach that performs Content-ID in a single unified engine and leverages a common signature format. This means that content is processed only once, and performance remains steady even as additional Content-ID features are enabled.

The single pass software uses a stream-based, uniform signature-matching engine for content inspection. Instead of using separate engines and signature sets (requiring multi-pass scanning) and instead of using file proxies (requiring file download prior to scanning), the single pass architecture scans traffic for all signatures once and in a stream-based fashion to avoid latency introduction.

The use of a stream-based engine replaces several components commonly used in other solutions—a file proxy for data, virus, and spyware, a signature engine for vulnerability exploits, and an http decoder for URL filtering. By using one common engine, two key benefits are realized. First, unlike file proxies that need to download the entire file before they can scan the traffic, a stream-based engine scans traffic real time, only reassembling packets as needed and only in very small amounts. Second, unlike traditional approaches, all traffic can be scanned with a single engine, instead of multiple scanning engines.

Content-ID is enabled on all Palo Alto Networks platforms through annual subscriptions for URL filtering and/or threat prevention, both of which provide support for unlimited users. The unlimited user support helps maintain a consistent annual cost structure while ensuring that new employees are protected as they are hired.

Palo Alto Networks single pass parallel processing architecture accelerates content inspection performance while minimizing latency.

Intrusion Prevention

Content-ID protects networks from all types of vulnerability exploits, buffer overflows, DoS attacks and port scans that lead to the compromise and damage of enterprise information resources. Palo Alto Networks IPS mechanisms include:

- Protocol decoders and anomaly detection
- Stateful pattern matching
- Statistical anomaly detection
- Heuristic-based analysis
- Invalid or malformed packet detection
- IP defragmentation and TCP reassembly
- Custom vulnerability and spyware phone-home signatures

Traffic is normalized to eliminate invalid and malformed packets, while TCP reassembly and IP de-fragmentation is performed to ensure the utmost accuracy and protection despite any packet-level evasion techniques.

WildFire: Protection from unknown malware and zero-day exploits

Criminals have increasingly turned to unknown malware and exploits to avoid traditional security controls. Palo Alto Networks has addressed this challenge with WildFire, which identifies unknown malware, zero-day exploits, and Advanced Persistent Threats (APTs) by observing their actual behavior in a virtualized environment, instead of relying solely on pre-existing signatures.

- **Integration of Firewall and the Cloud:** To support dynamic malware analysis across the network at scale, WildFire is built on a cloud-based architecture that can be leveraged by your existing Palo Alto Networks next-generation firewall, with no additional hardware. The in-line firewall captures unknown files and performs enforcement while maintaining high network throughput and low latency.
- **WildFire Virtualized Sandbox:** WildFire is an advanced, virtual malware analysis environment, purpose-built for high fidelity hardware emulation, analyzing suspicous samples as they execute. The cloud-based service detects and blocks targeted and unknown malware, exploits, and outbound C2 activity by observing their actual behavior, rather than relying on pre-existing signatures.
- **Automated Signature Generator:** When a sample is identified as malicious, WildFire automatically generates protections and delivers them to all WildFire customers globally in as little as 30 minutes.
- **Deep Visibility and Analysis:** WildFire users receive integrated logs, analysis, and visibility into WildFire events in the Palo Alto Networks management interface, Panorama, or the WildFire portal, enabling teams to quickly investigate and correlate events observed in their networks. This allows security staff to quickly locate the data needed for timely investigations and incident response. Host-based and network-based indicators of compromise become actionable through log analysis and custom signatures. An open API allows for integration with best-in-class SEIM tools, such as the Palo Alto Networks application for Splunk, and leading endpoint agents.

URL Filtering

Complementing the threat prevention and application control capabilities is a fully integrated, on-box URL filtering database that enables security teams to not only control end-user web surfing activities, but also combine URL context with application and user rules. The on-box URL database can be augmented to suit the traffic patterns of the local user community with a custom, 1 million URL database. URLs that are not categorized by the local URL database can be pulled into cache from a hosted, 180 million URL database. In addition to database customization, administrators can create custom URL categories to further tailor the URL controls to suit their specific needs. URL categorization can be combined with application and user classification to further target and define policies. For example, SSL decryption can be invoked for select high-risk URL categories to ensure threats are exposed, QoS controls can be applied to streaming media sites, URL filtering visibility and policy controls can be tied to specific users through the transparent integration with enterprise directory services (Active Directory, LDAP, eDirectory) with additional insight provided through customizable reporting and logging.

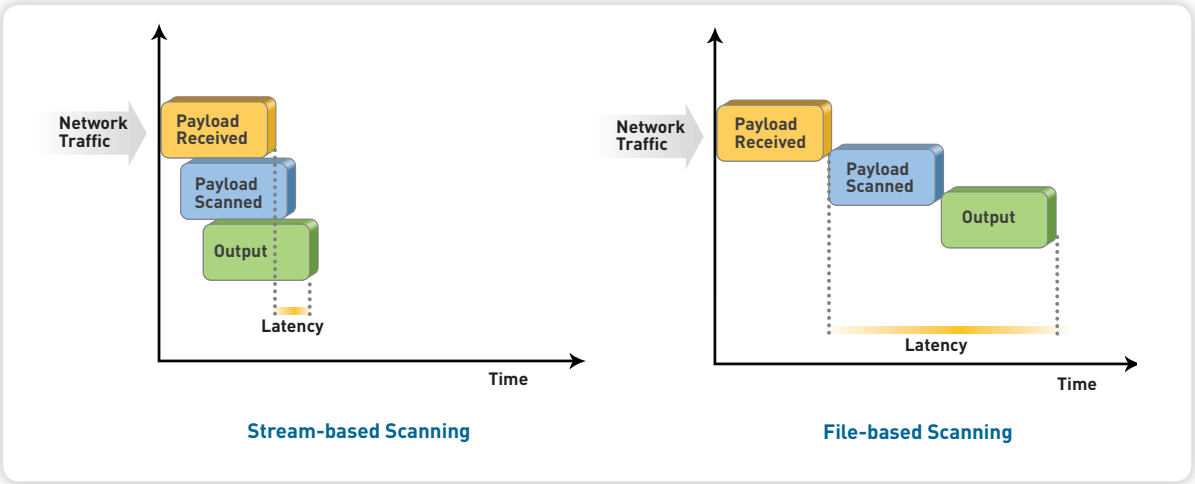
Administrators can configure a custom block page to notify end users of any policy violations. The page can include references to the username, IP address, the URL they are attempting to access and the URL category. In order to place some of the web activity ownership back in the user's hands, administrators can allow users to continue after being presented with a warning page, or can use passwords to override the URL filtering policy.

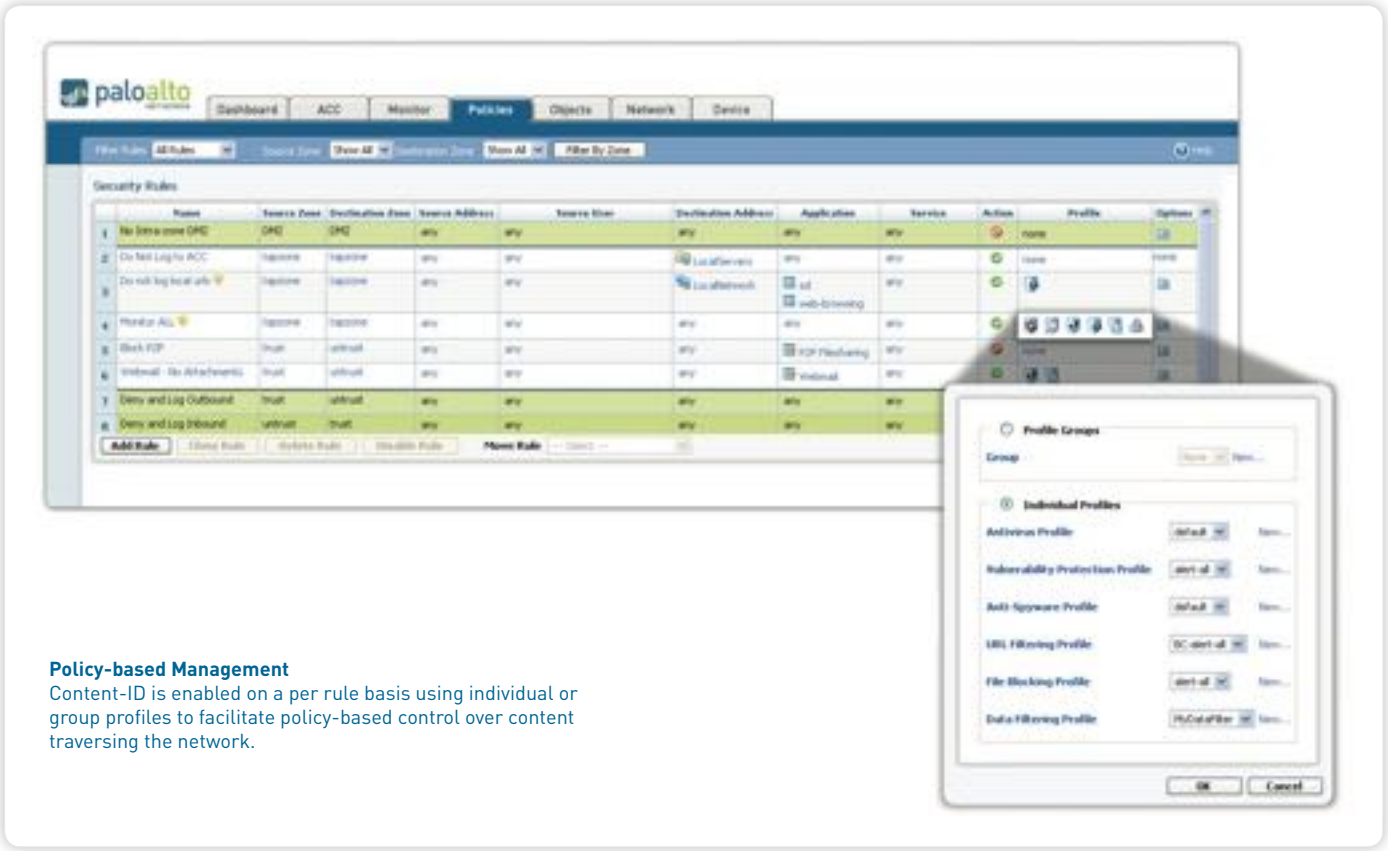
File and Data Filtering

Data filtering features enable administrators to implement policies that will reduce the risks associated with the transfer of unauthorized files and data.

- File blocking by type: Control the flow of a wide range of file types by looking deep within the payload to identify the file type (as opposed to looking only at the file extension).

Stream-based scanning
Stream-based scanning helps minimize latency and maximize throughput performance.





Policy-based Management
Content-ID is enabled on a per rule basis using individual or group profiles to facilitate policy-based control over content traversing the network.

- Data filtering: Control the transfer of sensitive data patterns such as credit card and social security numbers in application content or attachments.
- File transfer function control: Control the file transfer functionality within an individual application, allowing application use yet preventing undesired inbound or outbound file transfer.

Log Correlation and Reporting

Powerful log filtering enables administrators to quickly investigate security incidents by correlating threats with applications and user identity. The log viewer enables an administrator to click on a cell value to immediately create a filter that can be narrowed down further by combining multiple criteria using an expression builder and additional log fields, even if they are not visible in the log viewer. To tie the user identity to the threat, the log viewer leverages the integration with enterprise directory services. Log results can be exported to a CSV file for offline archival or further analysis. The trace session tool accelerates forensics and incident investigation with a centralized, correlated view across all of the logs for traffic, threats, URLs, and applications related to an individual session.

Reporting is enabled through a set of predefined reports that can be customized, pulling data from any of the log databases and then saving them for future use. Once the desired report is created, it can be configured to run on a regular basis, emailing a set of PDF reports or exporting them to CSV or PDF.

paloalto networks
the enterprise security company™

4401 Great America Parkway
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

Copyright ©2014, Palo Alto Networks, Inc. All rights reserved. Palo Alto Networks, the Palo Alto Networks Logo, PAN-OS, App-ID and Panorama are trademarks of Palo Alto Networks, Inc. All specifications are subject to change without notice. Palo Alto Networks assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document. Palo Alto Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. **PAN_TB_ContentID_093014**

GlobalProtect Administrator's Guide

8.0

paloaltonetworks.com/documentation

techDOCS paloalto NETWORKS

Contact Information

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

About the Documentation

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

Copyright

Palo Alto Networks, Inc.
www.paloaltonetworks.com

© 2017-2017 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

Last Revised

November 10, 2017

Table of Contents

GlobalProtect Overview 7
About the GlobalProtect Components 9
GlobalProtect Portal 9
GlobalProtect Gateways 9
GlobalProtect Client 9
What Client OS Versions are Supported with GlobalProtect? 11
About GlobalProtect Licenses 12

Get Started 13
Create Interfaces and Zones for GlobalProtect 15
Enable SSL Between GlobalProtect Components 17
About GlobalProtect Certificate Deployment 17
GlobalProtect Certificate Best Practices 17
Deploy Server Certificates to the GlobalProtect Components 20

Authentication 25
About GlobalProtect User Authentication 27
Supported GlobalProtect Authentication Methods 27
How Does the Agent or App Know What Credentials to Supply? 29
How Does the Agent Know Which Certificate to Supply? 30
Set Up External Authentication 31
Set Up LDAP Authentication 31
Set Up SAML Authentication 33
Set Up Kerberos Authentication 34
Set Up RADIUS or TACACS+ Authentication 35
Set Up Client Certificate Authentication 38
Deploy Shared Client Certificates for Authentication 38
Deploy Machine Certificates for Authentication 38
Deploy User-Specific Client Certificates for Authentication 42
Set Up Two-Factor Authentication 45
Enable Two-Factor Authentication Using Certificate and Authentication Profiles 45
Enable Two-Factor Authentication Using One-Time Passwords (OTPs) 47
Enable Two-Factor Authentication Using Smart Cards 50
Set Up Authentication for strongSwan Ubuntu and CentOS Clients 53
Enable Authentication Using a Certificate Profile 53
Enable Authentication Using an Authentication Profile 55
Enable Authentication Using Two-Factor Authentication 57
Configure GlobalProtect to Facilitate Multi-Factor Authentication Notifications 59
Enable Delivery of GlobalProtect Client VSAs to a RADIUS Server 62
Enable Group Mapping 63

GlobalProtect Gateways 65
GlobalProtect Gateways Overview 67
GlobalProtect Gateway Concepts 68
Types of Gateways 68
Gateway Priority in a Multiple Gateway Configuration 68

GlobalProtect MIB Support ... 69
Prerequisite Tasks for Configuring the GlobalProtect Gateway ... 71
Configure a GlobalProtect Gateway ... 72

GlobalProtect Portals ... 79
GlobalProtect Portal Overview ... 81
Prerequisite Tasks for Configuring the GlobalProtect Portal ... 82
Set Up Access to the GlobalProtect Portal ... 83
Define the GlobalProtect Client Authentication Configurations ... 85
Define the GlobalProtect Agent Configurations ... 86
Customize the GlobalProtect Agent ... 91
Customize the GlobalProtect Portal Login, Welcome, and Help Pages ... 99

GlobalProtect Clients ... 101
Deploy the GlobalProtect Client Software ... 103
Deploy the GlobalProtect Agent Software ... 103
Download and Install the GlobalProtect Mobile App ... 109
Download and Install the GlobalProtect App for Chrome OS ... 111
Deploy Agent Settings Transparently ... 115
Customizable Agent Settings ... 115
Deploy Agent Settings to Windows Clients ... 122
Deploy Agent Settings to Mac Clients ... 131

GlobalProtect Clientless VPN ... 135
Clientless VPN Overview ... 137
Supported Technologies ... 139
Configure Clientless VPN ... 140
Troubleshoot Clientless VPN ... 145

Mobile Endpoint Management ... 149
Mobile Endpoint Management Overview ... 151
Set Up a Mobile Endpoint Management System ... 152
Manage the GlobalProtect App Using AirWatch ... 153
Deploy the GlobalProtect Mobile App Using AirWatch ... 153
Configure the GlobalProtect App for iOS Using AirWatch ... 154
Configure the GlobalProtect App for Android Using AirWatch ... 157
Configure the GlobalProtect App for Windows 10 UWP Using AirWatch ... 161
Manage the GlobalProtect App Using a Third-Party MDM ... 164
Configure the GlobalProtect App for iOS ... 164
Configure the GlobalProtect App for Android ... 168

Host Information ... 171
About Host Information ... 173
What Data Does the GlobalProtect Agent Collect? ... 173
How Does the Gateway Use the Host Information to Enforce Policy? ... 175
How Do Users Know if Their Systems are Compliant? ... 176
How Do I Get Visibility into the State of the End Clients? ... 176
Configure HIP-Based Policy Enforcement ... 177
Collect Application and Process Data From Clients ... 184
Block Device Access ... 190

Configure Windows User-ID Agent to Collect Host Information 192
MDM Integration Overview 192
Information Collected 192
System Requirements 194
Configure GlobalProtect to Retrieve Host Information 194
Troubleshoot the MDM Integration Service 197

GlobalProtect Quick Configs 199
Remote Access VPN (Authentication Profile) 201
Remote Access VPN (Certificate Profile) 204
Remote Access VPN with Two-Factor Authentication 207
Always On VPN Configuration 211
Remote Access VPN with Pre-Logon 212
GlobalProtect Multiple Gateway Configuration 218
GlobalProtect for Internal HIP Checking and User-Based Access 221
Mixed Internal and External Gateway Configuration 225

GlobalProtect Architecture 231
GlobalProtect Reference Architecture Topology 233
GlobalProtect Portal 233
GlobalProtect Gateways 233
GlobalProtect Reference Architecture Features 235
End User Experience 235
Management and Logging 235
Monitoring and High Availability 236
GlobalProtect Reference Architecture Configurations 237
Gateway Configuration 237
Portal Configuration 237
Policy Configurations 237

GlobalProtect Cryptography 239
About GlobalProtect Cipher Selection 241
Cipher Exchange Between the GlobalProtect Agent and Gateway 242
GlobalProtect Cryptography References 244
Reference: GlobalProtect Agent Cryptographic Functions 244
TLS Cipher Suites Supported by GlobalProtect Agents 245
Ciphers Used to Set Up IPSec Tunnels 251
SSL APIs 253

GlobalProtect Overview

Whether checking email from home or updating corporate documents from the airport, the majority of today's employees work outside the physical corporate boundaries. This increased workforce mobility brings increased productivity and flexibility while simultaneously introducing significant security risks. Every time users leave the building with their laptops or mobile devices they are bypassing the corporate firewall and associated policies that are designed to protect both the user and the network. GlobalProtect™ solves the security challenges introduced by roaming users by extending the same next-generation firewall-based policies that are enforced within the physical perimeter to all users, no matter where they are located.

The following sections provide conceptual information about the Palo Alto Networks GlobalProtect offering and describe the components of GlobalProtect and the various deployment scenarios:

- About the GlobalProtect Components
- What Client OS Versions are Supported with GlobalProtect?
- What Features Does GlobalProtect Support?
- About GlobalProtect Licenses

© 2017 Palo Alto Networks, Inc.

About the GlobalProtect Components

GlobalProtect provides a complete infrastructure for managing your mobile workforce to enable secure access for all your users, regardless of what devices they are using or where they are located. This infrastructure includes the following components:

- GlobalProtect Portal
- GlobalProtect Gateways
- GlobalProtect Client

GlobalProtect Portal

The GlobalProtect portal provides the management functions for your GlobalProtect infrastructure. Every client system that participates in the GlobalProtect network receives configuration information from the portal, including information about available gateways as well as any client certificates that may be required to connect to the GlobalProtect gateway(s). In addition, the portal controls the behavior and distribution of the GlobalProtect agent software to both Mac and Windows laptops. (On mobile devices, the GlobalProtect app is distributed through the Apple App Store for iOS devices or through Google Play for Android devices.) If you are using the Host Information Profile (HIP) feature, the portal also defines what information to collect from the host, including any custom information you require. You Set Up Access to the GlobalProtect Portal on an interface on any Palo Alto Networks next-generation firewall.

GlobalProtect Gateways

GlobalProtect gateways provide security enforcement for traffic from GlobalProtect agents/apps. Additionally, if the HIP feature is enabled, the gateway generates a HIP report from the raw host data the clients submit and can use this information in policy enforcement. You can configure different Types of Gateways to provide security enforcement and/or virtual private network (VPN) access for your remote users, or to apply security policy for access to internal resources.

You Configure a GlobalProtect Gateway on an interface on any Palo Alto Networks next-generation firewall. You can run both a gateway and a portal on the same firewall, or you can have multiple, distributed gateways throughout your enterprise.

GlobalProtect Client

The GlobalProtect client software runs on end user systems and enables access to your network resources via the GlobalProtect portals and gateways you have deployed. There are two types of GlobalProtect clients:

- **The GlobalProtect Agent**—Runs on Windows and Mac OS systems and is deployed from the GlobalProtect portal. You configure the behavior of the agent—for example, which tabs the users can see, whether or not users can uninstall the agent—in the client configuration(s) you define on the portal. See Define the GlobalProtect Agent Configurations, Customize the GlobalProtect Agent, and Deploy the GlobalProtect Agent Software for details.
- **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).

See What Client OS Versions are Supported with GlobalProtect? for more details.

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

© 2017 Palo Alto Networks, Inc.



© 2017 Palo Alto Networks, Inc.

What Client OS Versions are Supported with GlobalProtect?

Palo Alto Networks supports the GlobalProtect app (also referred to as the GlobalProtect agent) on common desktop, laptop, and mobile devices. We recommend that you configure GlobalProtect on firewalls running PAN-OS 6.1 or a later release and that you install only supported releases of the GlobalProtect app on endpoints. The minimum GlobalProtect app release varies by operating system; to determine the minimum GlobalProtect app release for a specific operating system, refer to the following topics in the Palo Alto Networks® Compatibility Matrix:

- Where Can I Install the GlobalProtect App?
- What X-Auth IPSec Clients are Supported?

Older versions of the GlobalProtect app (releases 1.0 through 2.1) are still supported on the operating systems and PAN-OS releases with which they were released. For minimum PAN-OS release support for GlobalProtect app 2.1 and older releases, refer to the GlobalProtect agent (app) release notes for your specific release on the Software Updates site.

© 2017 Palo Alto Networks, Inc.

About GlobalProtect Licenses

If you simply want to use GlobalProtect to provide a secure, remote access or virtual private network (VPN) solution via single or multiple internal/external gateways, you do not need any GlobalProtect licenses. However, to use some of the more advanced features (such as HIP checks and associated content updates, support for the GlobalProtect mobile app, or IPv6 support) you need to purchase an annual GlobalProtect subscription. This license must be installed on each firewall running a gateway(s) that:

- Performs HIP checks
- Supports the GlobalProtect app on mobile devices
- Provides IPv6 connections

For GlobalProtect Clientless VPN, this feature also requires you to install a GlobalProtect subscription on the firewall that hosts the Clientless VPN from the GlobalProtect portal. You also need the **GlobalProtect Clientless VPN** dynamic updates to use this feature.

Feature	Subscription Required?
Single, external gateway (Windows and Mac)	
Single or multiple internal gateways	
Multiple external gateways	
HIP Checks	✓
Mobile app for iOS endpoints, Android endpoints, Chromebooks, and Windows 10 UWP endpoints	✓
IPv6 support	✓
Clientless VPN	✓

See Activate Licenses for information on installing licenses on the firewall.

Get Started

For GlobalProtect™ to work, you must set up the infrastructure that allows all of the components to communicate. At a basic level, this means setting up the interfaces and zones to which the GlobalProtect end users connect to access the portal and the gateways to the network. Because the GlobalProtect components communicate over secure channels, you must acquire and deploy the required SSL certificates to the various components. The following sections guide you through the steps to set up the GlobalProtect infrastructure:

> Create Interfaces and Zones for GlobalProtect on page 15
> Enable SSL Between GlobalProtect Components on page 17

© 2017 Palo Alto Networks, Inc.

Create Interfaces and Zones for GlobalProtect

You must configure the following interfaces and zones for your GlobalProtect infrastructure:

- **GlobalProtect portal**—Requires a Layer 3 or loopback interface for the GlobalProtect clients' connection. If the portal and gateway are on the same firewall, they can use the same interface. The portal must be in a zone that is accessible from outside your network, for example: DMZ.
- **GlobalProtect gateways**—The interface and zone requirements for the gateway depend on whether the gateway you are configuring is external or internal, as follows:
- **External gateways**—Requires a Layer 3 or loopback interface and a logical tunnel interface for the client to establish a VPN tunnel. The Layer 3/loopback interface must be in an external zone, such as DMZ. A tunnel interface can be in the same zone as the interface connecting to your internal resources (for example trust). For added security and better visibility, you can create a separate zone, such as corp-vpn. If you create a separate zone for your tunnel interface, you must create security policies that enable traffic to flow between the VPN zone and the trust zone.
- **Internal gateways**—Requires a Layer 3 or loopback interface in your trust zone. You can also create a tunnel interface for access to your internal gateways, but this is not required.



For tips on how to use a loopback interface to provide access to GlobalProtect on different ports and addresses, refer to Can GlobalProtect Portal Page be Configured to be Accessed on any Port?

For more information about portals and gateways, see About the GlobalProtect Components.

STEP 1 | Configure a Layer 3 interface for each portal and/or gateway you plan to deploy.



If the gateway and portal are on the same firewall, you can use a single interface for both.



As a best practice use static IP addresses for the portal and gateway.



Do not attach an interface management profile that allows HTTP, HTTPS, Telnet, or SSH on the interface where you have configured a GlobalProtect portal or gateway because this enables access to your management interface from the Internet. Follow the Best Practices for Securing Administrative Access *to ensure that you are securing administrative access to your firewalls in a way that will prevent successful attacks.*

1. Select **Network** > **Interfaces** > **Ethernet** or **Network** > **Interfaces** > **Loopback** and then select the interface you want to configure for GlobalProtect. In this example, we are configuring ethernet1/1 as the portal interface.
2. (**Ethernet only**) Select **Layer3** from the **Interface Type** drop-down.
3. On the **Config** tab, select the zone to which the portal or gateway interface belongs as follows:
   - Place portals and external gateways in an untrust zone for access by hosts outside your network, such as l3-untrust.
   - Place internal gateways in an internal zone, such as l3-trust.
   - If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for the new zone and then click **OK**.
4. In the **Virtual Router** drop-down, select **default**.
5. Assign an IP address to the interface:

© 2017 Palo Alto Networks, Inc.

- For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
- For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.

6. To save the interface configuration, click **OK**.

STEP 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect agents.

IP addresses are not required on the tunnel interface unless you require dynamic routing. In addition, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.

Be sure to enable User-ID in the zone where the VPN tunnels terminate.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
   - (**Recommended**) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example, corp-vpn), select the **Enable User Identification** check box, and then click **OK**.
4. In the **Virtual Router** drop-down, select **None**.
5. Assign an IP address to the interface:
   - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
   - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
6. To save the interface configuration, click **OK**.

STEP 3 | If you created a separate zone for tunnel termination of VPN connections, create a security policy to enable traffic flow between the VPN zone and your trust zone.

For example, the following policy rule enables traffic between the corp-vpn zone and the l3-trust zone.



STEP 4 | Save the configuration.

© 2017 Palo Alto Networks, Inc.

Enable SSL Between GlobalProtect Components

All interaction between the GlobalProtect components occurs over an SSL/TLS connection. Therefore, you must generate and/or install the required certificates before configuring each component so that you can reference the appropriate certificate(s) in the configurations. The following sections describe the supported methods of certificate deployment, descriptions and best practice guidelines for the various GlobalProtect certificates, and provide instructions for generating and deploying the required certificates:

- About GlobalProtect Certificate Deployment on page 17
- GlobalProtect Certificate Best Practices on page 17
- Deploy Server Certificates to the GlobalProtect Components on page 20

About GlobalProtect Certificate Deployment

There are three basic approaches to Deploy Server Certificates to the GlobalProtect Components on page 20:

- **(Recommended) Combination of third-party certificates and self-signed certificates**—Because the end clients will be accessing the portal prior to GlobalProtect configuration, the client must trust the certificate to establish an HTTPS connection.
- **Enterprise Certificate Authority**—If you already have your own enterprise CA, you can use this internal CA to issue certificates for each of the GlobalProtect components and then import them onto the firewalls hosting your portal and gateway(s). In this case, you must also ensure that the end user systems/mobile devices trust the root CA certificate used to issue the certificates for the GlobalProtect services to which they must connect.
- **Self-Signed Certificates**—You can generate a self-signed CA certificate on the portal and use it to issue certificates for all of the GlobalProtect components. However, this solution is less secure than the other options and is therefore not recommended. If you do choose this option, end users will see a certificate error the first time they connect to the portal. To prevent this, you can deploy the self-signed root CA certificate to all end user systems manually or using some sort of centralized deployment, such as an Active Directory Group Policy Object (GPO).

GlobalProtect Certificate Best Practices

The following table summarizes the SSL/TLS certificates you will need, depending on which features you plan to use:

Certificate	Usage	Issuing Process/Best Practices
CA certificate	Used to sign certificates issued to the GlobalProtect components.	If you plan to use self-signed certificates, a best practice is to generate a CA certificate on the portal and then use that certificate to issue the required GlobalProtect certificates.
Portal server certificate	Enables GlobalProtect agents and apps to establish an HTTPS connection with the portal.	• This certificate is identified in an SSL/TLS service profile. You assign the portal server certificate by selecting its associated service profile in a portal configuration. • Use a certificate from a well-known, third-party CA. This is the most secure option and ensures

© 2017 Palo Alto Networks, Inc.

Certificate	Usage	Issuing Process/Best Practices
		that the user endpoints can establish a trust relationship with the portal and without requiring you to deploy the root CA certificate. • If you do not use a well-known, public CA, you should export the root CA certificate that was used to generate the portal server certificate to all endpoints that run the GlobalProtect agent or application. Exporting this certificate prevents the end users from seeing certificate warnings during the initial portal login. • The Common Name (CN) and, if applicable, the Subject Alternative Name (SAN) fields of the certificate must match the IP address or FQDN of the interface that hosts the portal. • In general, a portal must have its own server certificate. However, if you are deploying a single gateway and portal on the same interface for basic VPN access, you must use the same certificate for both the gateway and the portal. • If you configure a gateway and portal on the same interface, we also recommend that you use the same certificate profile and SSL/TLS service profile for both the gateway and portal. If they do not use the same certificate profile and SSL/TLS service profile, the gateway configuration takes precedence over the portal configuration during the SSL handshake.
Gateway server certificate	Enables GlobalProtect agents and apps to establish an HTTPS connection with the gateway.	• This certificate is identified in an SSL/TLS service profile. You assign the portal server certificate by selecting its associated service profile in a gateway configuration. • Generate a CA certificate on the portal and use that CA certificate to generate all gateway certificates. • The CN and, if applicable, the SAN fields of the certificate must match the FQDN or IP address of the interface where you plan to configure the gateway. • The portal distributes the gateway root CA certificates to agents in the client configuration, so the gateway certificates do not need to be issued by a public CA. • In general, each gateway must have its own server certificate. However, if you are deploying a single gateway and portal on the same interface for basic VPN access, you must use a single server certificate for both components.As a best practice, use a certificate that a public CA signed. • If you configure a gateway and portal on the same interface, we also recommend that you use the same certificate profile and SSL/TLS service

© 2017 Palo Alto Networks, Inc.

Certificate	Usage	Issuing Process/Best Practices
		profile for both the gateway and portal. If they do not use the same certificate profile and SSL/TLS service profile, the gateway configuration takes precedence over the portal configuration during the SSL handshake.

(**Optional**) Client certificate	Used to enable mutual authentication in establishing an HTTPS session between the GlobalProtect agents and the gateways/portal. This ensures that only devices with valid client certificates are able to authenticate and connect to the network.	• For simplified deployment of client certificates, configure the portal to deploy the client certificate to the agents upon successful login using either of the following methods:
		  • Use a single client certificate across all GlobalProtect agents that receive the same configuration. You assign the **Local** client certificate by uploading the certificate to the portal and selecting it in a portal agent configuration.
		  • Use simple certificate enrollment protocol (**SCEP**) to enable the GlobalProtect portal to deploy unique client certificates to your GlobalProtect agents. You enable this by configuring a SCEP profile and then selecting that profile in a portal agent configuration.
		• Use one of the following supported digest algorithms when you generate client certificates for GlobalProtect endpoints: sha1, sha256, or sha384. Sha512 is not supported with client certificates.
		• You can use other mechanisms to deploy unique client certificates to each client system for use in authenticating the end user.
		• Consider testing your configuration without the client certificate first, and then add the client certificate after you are sure that all other configuration settings are correct.

(**Optional**) Machine certificates	A machine certificate is a client certificate that is issued to a device. Each machine certificate identifies the device in the subject field (for example, CN=laptop1.example.com) instead of a user. The certificate ensures that only trusted endpoints can connect to gateways or the portal. Machine certificates are required for users whose connect method is pre-logon, which enables	• Use one of the following supported digest algorithms when you generate client certificates for GlobalProtect endpoints: sha1, sha256, or sha384. Sha512 is not supported with client certificates. • If you plan to use the pre-logon feature, use your own PKI infrastructure to deploy machine certificates to each client system prior to enabling GlobalProtect access. This approach is important for ensuring security. For more information, see Remote Access VPN with Pre-Logon.

© 2017 Palo Alto Networks, Inc.

Certificate	Usage	Issuing Process/Best Practices
	GlobalProtect to establish a VPN tunnel before the user logs in.	

Table: GlobalProtect Certificate Requirements

For details about the types of keys for secure communication between the GlobalProtect endpoint and the portals and gateways, see Reference: GlobalProtect Agent Cryptographic Functions.

Deploy Server Certificates to the GlobalProtect Components

The following table shows the best practice steps for deploying SSL/TLS certificates to the GlobalProtect components:

- Import a server certificate from a well-known, third-party CA.

  Use a server certificate from a well-known, third-party CA for the GlobalProtect portal. This practice ensures that the end users are able to establish an HTTPS connection without seeing warnings about untrusted certificates.

  The CN and, if applicable, the SAN fields of the certificate must match the FQDN or IP address of the interface where you plan to configure the portal or the device check-in interface on a third-party mobile endpoint management system. Wildcard matches are supported.

  Before you import a certificate, make sure the certificate and key files are accessible from your management system and that you have the passphrase to decrypt the private key.

  1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
  2. Click **Import**.
  3. Use the **Local** certificate type (the default).
  4. Enter a **Certificate Name**.
  5. Enter the path and name to the **Certificate File** received from the CA, or **Browse** to find the file.
  6. Select **Encrypted Private Key and Certificate (PKCS12)** as the **File Format**.
  7. Enter the path and name to the PKCS#12 file in the **Key File** field or **Browse** to find it.
  8. Enter and re-enter the **Passphrase** that was used to encrypt the private key and then click **OK** to import the certificate and key.

- Create the root CA certificate for issuing self-signed certificates for the GlobalProtect components.

  Create the Root CA certificate on the portal and use it to issue server certificates for the gateways and, optionally, for clients.

  Before deploying self-signed certificates, you must create the root CA certificate that signs the certificates for the GlobalProtect components:

  1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**.
  2. Use the **Local** certificate type (the default).
  3. Enter a **Certificate Name**, such as GlobalProtect_CA. The certificate name cannot contain spaces.

© 2017 Palo Alto Networks, Inc.

4. Do not select a value in the **Signed By** field. (Without a selection for **Signed By**, the certificate is self-signed.)
5. Select the **Certificate Authority** check box.
6. Click **OK** to generate the certificate.

- Use the root CA on the portal to generate a self-signed server certificate.



Generate server certificates for each gateway you plan to deploy and optionally for the management interface of the third-party mobile endpoint management system (if this interface is where the gateways retrieve HIP reports).



In the gateway server certificates, the values in the CN and SAN fields must be identical. If the values differ, the GlobalProtect agent detects the mismatch and does not trust the certificate. Self-signed certificates contain a SAN field only if you add a Host Name attribute.

As an alternative method, you can Use Simple Certificate Enrollment Protocol (SCEP) to request a server certificate from your enterprise CA.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**.
2. Use the **Local** certificate type (the default).
3. Enter a **Certificate Name**. This name cannot contain spaces.
4. In the **Common Name** field, enter the FQDN (**recommended**) or IP address of the interface where you plan to configure the gateway.
5. In the **Signed By** field, select the GlobalProtect_CA you created.
6. In the Certificate Attributes section, **Add** and define the attributes that uniquely identify the gateway. Keep in mind that if you add a **Host Name** attribute (which populates the SAN field of the certificate), it must be the same as the value you defined for the **Common Name**.
7. Configure cryptographic settings for the server certificate including encryption **Algorithm**, key length (**Number of Bits**), **Digest** algorithm and **Expiration** (days).
8. Click **OK** to generate the certificate.

- Use Simple Certificate Enrollment Protocol (SCEP) to request a server certificate from your enterprise CA.



Configure separate SCEP profiles for each portal and gateway you plan to deploy. Then use the specific SCEP profile to generate the server certificate for each GlobalProtect component.



*In portal and gateway server certificates, the value of the CN field must include the FQDN (**recommended**) or IP address of the interface where you plan to configure the portal or gateway and must be identical to the SAN field.*



To comply with the U.S. Federal Information Processing Standard (FIPS), you must also enable mutual SSL authentication between the SCEP server and the GlobalProtect portal. (FIPS-CC operation is indicated on the firewall login page and in its status bar.)

After you commit the configuration, the portal attempts to request a CA certificate using the settings in the SCEP profile. If successful, the firewall hosting the portal saves the CA certificate and displays it in the list of **Device Certificates**.

1. Configure a SCEP Profile for each GlobalProtect portal or gateway:

© 2017 Palo Alto Networks, Inc.

1. Enter a **Name** that identifies the SCEP profile and the component to which you deploy the server certificate. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the **Location** where the profile is available.
2. (**Optional**) Configure a **SCEP Challenge**-response mechanism between the PKI and portal for each certificate request. Use either a **Fixed** challenge password which you obtain from the SCEP server or a **Dynamic** password where the portal-client submits a username and OTP of your choice to the SCEP Server. For a Dynamic SCEP challenge, this can be the credentials of the PKI administrator.
3. Configure the **Server URL** that the portal uses to reach the SCEP server in the PKI (for example, **`http://10.200.101.1/certsrv/mscep/`**).
4. Enter a string (up to 255 characters in length) in the **CA-IDENT Name** field to identify the SCEP server.
5. Enter the **Subject** name to use in the certificates generated by the SCEP server. The subject must include a common name (CN) key in the format **`CN=<value>`** where **`<value>`** is the FQDN or IP address of the portal or gateway.
6. Select the **Subject Alternative Name Type**. To enter the email name in a certificate's subject or Subject Alternative Name extension, select **RFC 822 Name**. You can also enter the **DNS Name** to use to evaluate certificates, or the **Uniform Resource Identifier** to identify the resource from which the client will obtain the certificate.
7. Configure additional cryptographic settings including the key length (**Number of Bits**), and **Digest** algorithm for the certificate signing request.
8. Configure the permitted uses of the certificate, either for signing (**Use as digital signature**) or encryption (**Use for key encipherment**).
9. To ensure that the portal is connecting to the correct SCEP server, enter the **CA Certificate Fingerprint**. Obtain this fingerprint from the SCEP server interface in the Thumbprint field.
10. Enable mutual SSL authentication between the SCEP server and the GlobalProtect portal.
11. Click **OK** and then **Commit** the configuration.

2. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**.
3. Enter a **Certificate Name**. This name cannot contain spaces.
4. Select the **SCEP Profile** to use to automate the process of issuing a server certificate that is signed by the enterprise CA to a portal or gateway, and then click **OK** to generate the certificate. The GlobalProtect portal uses the settings in the SCEP profile to submit a CSR to your enterprise PKI.

- Assign the server certificate you imported or generated to an SSL/TLS service profile.
  1. Select **Device** > **Certificate Management** > **SSL/TLS Service Profile** and click **Add**.
  2. Enter a **Name** to identify the profile and select the server **Certificate** you imported or generated.
  3. Define the range of SSL/TLS versions (**Min Version** to **Max Version**) for communication between GlobalProtect components.

     To provide the strongest security, set the Min Version to TLSv1.2.
  4. Click **OK** to save the SSL/TLS service profile.
  5. **Commit** the changes.

- Deploy the self-signed server certificates.

  - *Export the self-signed server certificates issued by the root CA on the portal and import them onto the gateways.*
  - *Be sure to issue a unique server certificate for each gateway.*

© 2017 Palo Alto Networks, Inc.

- *If specifying self-signed certificates, you must distribute the Root CA certificate to the end clients in the portal client configurations.*

Export the certificate from the portal:

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
2. Select the gateway certificate you want to deploy and click **Export**.
3. In the **File Format** drop-down, select **Encrypted Private Key and Certificate (PKCS12)**.
4. Enter (and re-enter) a **Passphrase** to encrypt the private key.
5. Click **OK** to download the PKCS12 file to a location of your choice.

Import the certificate on the gateway:

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
2. Click **Import**.
3. Enter a **Certificate Name**.
4. **Browse** to find and select the **Certificate File** you downloaded in the previous step.
5. In the **File Format** drop-down, select **Encrypted Private Key and Certificate (PKCS12)**.
6. Enter (and re-enter) the **Passphrase** you used to encrypt the private key when you exported it from the portal.
7. Click **OK** to import the certificate and key.
8. **Commit** the changes to the gateway.

© 2017 Palo Alto Networks, Inc.

Authentication

The GlobalProtect™ portal and gateway must authenticate the end-user before it allows access to GlobalProtect resources. You must configure authentication mechanisms before continuing with the portal and gateway setup. The following sections detail the supported authentication mechanisms and how to configure them:

> About GlobalProtect User Authentication
> Set Up External Authentication
> Set Up Client Certificate Authentication
> Set Up Two-Factor Authentication
> Set Up Authentication for strongSwan Ubuntu and CentOS Clients
> Configure GlobalProtect to Facilitate Multi-Factor Authentication Notifications
> Enable Delivery of GlobalProtect Client VSAs to a RADIUS Server
> Enable Group Mapping

© 2017 Palo Alto Networks, Inc.

About GlobalProtect User Authentication

The first time a GlobalProtect client connects to the portal, the user is prompted to authenticate to the portal. If authentication succeeds, the GlobalProtect portal sends the GlobalProtect configuration, which includes the list of gateways to which the agent can connect, and optionally a client certificate for connecting to the gateways. After successfully downloading and caching the configuration, the client attempts to connect to one of the gateways specified in the configuration. Because these components provide access to your network resources and settings, they also require the end user to authenticate.

The appropriate level of security required on the portal and gateways varies with the sensitivity of the resources that the gateway protects. GlobalProtect provides a flexible authentication framework that allows you to choose the authentication profile and certificate profile that are appropriate to each component.

- Supported GlobalProtect Authentication Methods
- How Does the Agent or App Know What Credentials to Supply?
- How Does the Agent Know Which Certificate to Supply?

Supported GlobalProtect Authentication Methods

The following topics describe the authentication methods that GlobalProtect supports and provide usage guidelines for each method.

- Local Authentication on page 27
- External Authentication on page 27
- Client Certificate Authentication on page 27
- Two-Factor Authentication on page 28
- Multi-Factor Authentication for Non-Browser-Based Applications on page 29

Local Authentication

Both the user account credentials and the authentication mechanisms are local to the firewall. This authentication mechanism is not scalable because it requires an account for every GlobalProtect user and is, therefore, advisable for only very small deployments.

External Authentication

The user authentication functions are performed by an external LDAP, Kerberos, TACACS+, SAML, or RADIUS service (including support for two-factor, token-based authentication mechanisms, such as one-time password (OTP) authentication). To enable external authentication:

- Create a server profile with settings for access to the external authentication service.
- Create an authentication profile that refers to the server profile.
- Specify client authentication in the portal and gateway configurations and optionally specify the OS of the endpoint that will use these settings.

You can use different authentication profiles for each GlobalProtect component. See Set Up External Authentication on page 31 for instructions. See Remote Access VPN (Authentication Profile) on page 201 for an example configuration.

Client Certificate Authentication

For enhanced security, you can configure the portal or gateway to use a client certificate to obtain the username and authenticate the user before granting access to the system.

© 2017 Palo Alto Networks, Inc.

- To authenticate the user, one of the certificate fields, such as the Subject Name field, must identify the username.
- To authenticate the endpoint, the Subject field of the certificate must identify the device type instead of the username. (With the pre-logon connect methods, the portal or gateway authenticates the endpoint before the user logs in.)

For an agent configuration profile that specifies client certificates, each user receives a client certificate. The mechanism for providing the certificates determines whether a certificate is unique to each client or the same for all clients under that agent configuration:

- To deploy client certificates that are unique to each user and device, use **SCEP**. When a user first logs in, the portal requests a certificate from the enterprise's PKI. The portal obtains a unique certificate and deploys it to the client.
- To deploy the same client certificate to all users that receive an agent configuration, deploy a certificate that is **Local** to the firewall.

Use an optional certificate profile to verify the client certificate that a client presents with a connection request. The certificate profile specifies the contents of the username and user domain fields; lists CA certificates; criteria for blocking a session; and offers ways to determine the revocation status of CA certificates. Because the certificate is part of the authentication of the endpoint or user for a new session, you must pre-deploy certificates used in certificate profiles to the endpoints before the users' initial portal login.

The certificate profile specifies which certificate field contains the username. If the certificate profile specifies Subject in the Username Field, the certificate presented by the client must contain a common-name for the client to connect. If the certificate profile specifies a Subject-Alt with an Email or Principal Name as the Username Field, the certificate from the client must contain the corresponding fields, which will be used as the username when the GlobalProtect agent authenticates to the portal or gateway.

GlobalProtect also supports authentication by common access cards (CACs) and smart cards, which rely on a certificate profile. With these cards, the certificate profile must contain the root CA certificate that issued the certificate to the smart card or CAC.

If you specify client certificate authentication, you should not configure a client certificate in the portal configuration because the client system provides it when the user connects. For an example of how to configure client certificate authentication, see Remote Access VPN (Certificate Profile) on page 204.

Two-Factor Authentication

With two-factor authentication, the portal or gateway uses two mechanisms to authenticate a user, such as a one-time password in addition to AD login credentials. You can enable two-factor authentication on the portal and gateways by configuring a certificate profile and an authentication profile and adding them both to the portal and/or gateway configuration.

You can configure the portal and gateways to use the same authentication methods or use different methods. Regardless, with two-factor authentication, the client must successfully authenticate by the two mechanisms that the component demands before it grants access.

If the certificate profile specifies a Username Field from which GlobalProtect can obtain a username, the external authentication service automatically uses the username to authenticate the user to the external authentication service specified in the authentication profile. For example, if the Username Field in the certificate profile is set to Subject, the value in the common-name field of the certificate is used as the username when the authentication server tries to authenticate the user. If you do not want to force users to authenticate with a username from the certificate, make sure the certificate profile is set to None for the Username Field. See Remote Access VPN with Two-Factor Authentication on page 207 for an example configuration.

© 2017 Palo Alto Networks, Inc.

Multi-Factor Authentication for Non-Browser-Based Applications

(Windows and Mac endpoints only) For sensitive, non-browser-based network resources (for example, financial applications or software development applications) that may require additional authentication, GlobalProtect clients can now notify and prompt the user to perform the timely, multi-factor authentication needed to access these resources.

How Does the Agent or App Know What Credentials to Supply?

By default, the GlobalProtect agent attempts to use the same login credentials for the gateway that it used for portal login. In the simplest case, where the gateway and the portal use the same authentication profile and/or certificate profile, the agent will connect to the gateway transparently.

On a per-agent configuration basis, you can also customize which GlobalProtect portal and gateways—internal, external, or manual only—require different credentials (such as unique OTPs). This enables the GlobalProtect portal or gateway to prompt for the unique OTP without first prompting for the credentials specified in the authentication profile.

There are two options for modifying the default agent authentication behavior so that authentication is both stronger and faster:

- Cookie Authentication on the Portal or Gateway on page 29
- Credential Forwarding to Some or All Gateways on page 29

Cookie Authentication on the Portal or Gateway

Cookie authentication simplifies the authentication process for end users because they will no longer be required to log in to both the portal and the gateway in succession or enter multiple OTPs for authenticating to each. This improves the user experience by minimizing the number of times that users must enter credentials. In addition, cookies enable use of a temporary password to re-enable VPN access after the user's password expires.

You can configure cookie authentication settings independently for the portal and for individual gateways, (for example, you can impose a shorter cookie lifetime on gateways that protect sensitive resources). After the portal or gateways deploy an authentication cookie to the endpoint, the portal and gateways both rely on the same cookie to authenticate the user. When the agent presents the cookie, the portal or gateway evaluates whether the cookie is valid based on the configured cookie lifetime. If the cookie expires, GlobalProtect automatically prompts the user to authenticate with the portal or gateway. When authentication is successful, the portal or gateway issues the replacement authentication cookie to the endpoint and the validity period starts over.

Consider the following example where you configure the cookie lifetime for the portal—which does not protect sensitive information—as 15 days, but configure the cookie lifetime for gateways—which do protect sensitive information—as 24 hours. When the user first authenticates with the portal, the portal issues the authentication cookie. If after five days, the user attempted to connect to the portal, the authentication cookie would still be valid. However, if after five days the user attempted to connect to the gateway, the gateway would evaluate the cookie lifetime and determine it expired (5 days > 24 hours). The agent would then automatically prompt the user to authenticate with the gateway and, on successful authentication, receive a replacement authentication cookie. The new authentication cookie would then be valid for another 15 days on the portal and another 24 hours on the gateways.

For an example of how to use this option, see Set Up Two-Factor Authentication on page 45.

Credential Forwarding to Some or All Gateways

With two-factor authentication, you can specify the portal and/or types of gateways (internal, external, or manual only) that prompt for their own set of credentials. This option speeds up the authentication process when the portal and the gateway require different credentials (either different OTPs or different

© 2017 Palo Alto Networks, Inc.

login credentials entirely). For each portal or gateway that you select, the agent will not forward credentials, allowing you to customize the security for different GlobalProtect components. For example, you can have the same security on your portals and internal gateways, while requiring a second factor OTP or a different password for access to those gateways that provide access to your most sensitive resources.

For an example of how to use this option, see Set Up Two-Factor Authentication on page 45.

How Does the Agent Know Which Certificate to Supply?

When you configure GlobalProtect to use client certificates for authentication on Mac or Windows endpoints, GlobalProtect must present a valid client certificate to authenticate with the portal and/or gateways.

For a client certificate to be valid, it must meet the following four requirements:

- Issued by the certificate authority (CA) you defined in the Certificate Profile of your portal and gateway configurations.
- Specifies the client authentication purpose, which the certificate administrator specifies when creating the certificate.
- Located in the certificate store as configured in the GlobalProtect portal agent configuration. By default, the GlobalProtect agent first looks for a valid certificate in the user store and, if none exists, then looks in the machine store. Because the user store takes precedence, if the GlobalProtect agent locates a certificate in the user store, it will not look in the machine store. To force the GlobalProtect agent to look for certificates in only one certificate store, configure the **Client Certificate Store Lookup** option in the appropriate GlobalProtect portal agent configuration.
- Match any additional purpose specified in the GlobalProtect portal agent configuration. To specify an additional purpose, you must identify the object identifier (OID) for the certificate and configure the **Extended Key Usage OID** value in the appropriate GlobalProtect portal agent configuration. An OID is a numeric value that identifies the application or service for which to use a certificate and that is automatically attached to a certificate when it is created by a certificate authority (CA). For more information on specifying a common or custom OID, see Certificate Selection by OID.

When only one client certificate meets the requirements above, the agent automatically selects and uses that client certificate for authentication. However, when multiple client certificates meet the these requirements, GlobalProtect prompts the user to select the client certificate from a list of valid client certificates on the endpoint. While GlobalProtect requires users to select the client certificate only when they first connect, users might not know which certificate to select. In this case, we recommend you to narrow the list of available client certificates by certificate purpose (as indicated by the OID) and certificate store. For more information on these and other settings you can configure to customize your agent, see Customize the GlobalProtect Agent.

© 2017 Palo Alto Networks, Inc.

Set Up External Authentication

The following workflows describe how to set up the GlobalProtect portal and gateways to use an external authentication service. The supported authentication services are LDAP, Kerberos, RADIUS, SAML, or TACACS+.

These workflows also describe how to create an optional *authentication profile* that a portal or gateway can use to identify the external authentication service. This step is optional for external authentication because the authentication profile also can specify the local authentication database or None.



GlobalProtect also supports local authentication. To use local authentication, create a local user database (Device > Local User Database) that contains the users and groups to which you want to allow VPN access and then refer to that database in the authentication profile.

For more information, see Supported GlobalProtect Authentication Methods on page 27 or watch a video.

The options for setting up external authentication include:

- Set Up LDAP Authentication on page 31
- Set Up SAML Authentication on page 33
- Set Up Kerberos Authentication on page 34
- Set Up RADIUS or TACACS+ Authentication on page 35

Set Up LDAP Authentication

LDAP is often used by organizations as a central repository for user information and as an authentication service. It can also be used to store the role information for application users.

STEP 1 | Create a server profile.

The server profile identifies the external authentication service and instructs the firewall how to connect to that authentication service and access the authentication credentials for your users.



When you use LDAP to connect to Active Directory (AD), you must create a separate LDAP server profile for every AD domain.

1. Select **Device** > **Server Profiles** and select the **LDAP** profile.
2. Click **Add** and enter a Profile **Name**, such as GP-User-Auth.
3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the Location where the profile is available.
4. Select the **Type** of LDAP server.
5. Click **Add** in the Servers section and then enter the necessary information for connecting to the authentication server, including the server **Name**, IP address or FQDN of the **Server**, and **Port**.
6. Specify settings to enable the authentication service to authenticate the firewall. Enter the **Bind DN** and **Password**.
7. If you want the device to use SSL or TLS for a more secure connection with the directory server, select the **Require SSL/TLS secured connection** check box (selected by default). The protocol that the device uses depends on the server **Port**:
   - 389 (**default**)—TLS (Specifically, the device uses the StartTLS operation, which upgrades the initial plaintext connection to TLS.)
   - 636—SSL
   - Any other port—The device first attempts to use TLS. If the directory server doesn't support TLS, the device falls back to SSL.

© 2017 Palo Alto Networks, Inc.

8. For additional security, select the **Verify Server Certificate for SSL sessions** check box so that the device verifies the certificate that the directory server presents for SSL/TLS connections. To enable verification, you also have to select the **Require SSL/TLS secured connection** check box. For verification to succeed, the certificate must meet one of the following conditions:
   - It is in the list of device certificates: **Device** > **Certificate Management** > **Certificates** > **Device Certificates.** Import the certificate into the device, if necessary.
   - The certificate signer is in the list of trusted certificate authorities: **Device** > **Certificate Management** > **Certificates** > **Default Trusted Certificate Authorities**.
9. Click **OK** to save the server profile.

STEP 2 | (**Optional**) Create an authentication profile.

The authentication profile specifies the server profile for the portal or gateways to use when they authenticate users. On a portal or gateway, you can assign one or more authentication profiles in one or more *client authentication* profiles. For descriptions of how an authentication profile within a client authentication profile supports granular user authentication, see Configure a GlobalProtect Gateway and Set Up Access to the GlobalProtect Portal.



To enable users to connect and change their own expired passwords without administrative intervention, consider using a pre-logon connect method. See Remote Access VPN with Pre-Logon *for details.*

If users allow their passwords to expire, you may assign a temporary LDAP password to enable them to log in to the VPN. In this case, the temporary password may be used to authenticate to the portal, but the gateway login may fail because the same temporary password cannot be re-used. To prevent this, enable an authentication override in the portal configuration (Network > GlobalProtect > Portal) to enable the agent to use a cookie to authenticate to the portal and use the temporary password to authenticate the gateway.

1. Select **Device** > **Authentication Profile** and **Add** a new profile.
2. Enter a **Name** for the profile and then select **LDAP** as the authentication **Type**.
3. Select the LDAP authentication **Server Profile** that you created in 1 from the drop-down.
4. Enter **sAMAccountName** as the **Login Attribute**.
5. Set the **Password Expiry Warning** to specify the number of days before password expiration that users will be notified. By default, users will be notified seven days prior to password expiration (range is 1-255). Because users must change their passwords before the end of the expiration period, make sure you provide a notification period that is adequate for your user base to ensure continued access to the VPN. To use this feature, you must specify one of the following types of LDAP servers in your LDAP server profile: **active-directory**, **e-directory**, or **sun**.

   Users cannot access the VPN if their passwords expire unless you enable pre-logon.
6. Configure an optional custom expiry message to include additional instructions, such as help desk contact information or a link to a password portal where users can change their passwords (see Customize the GlobalProtect Agent).
7. Specify the domain name and username format. The device combines the **User Domain** and **Username Modifier** values to modify the domain/username string that a user enters during login. The device uses the modified string for authentication and uses the **User Domain** value for User-ID group mapping. Modifying user input is useful when the authentication service requires domain/username strings in a particular format and you don't want to rely on users to correctly enter the domain. You can select from the following options:
   - To send only the unmodified user input, leave the **User Domain** blank (the default) and set the **Username Modifier** to the variable **%USERINPUT%** (the default).

© 2017 Palo Alto Networks, Inc.

- To prepend a domain to the user input, enter a **User Domain** and set the **Username Modifier** to **%USERDOMAIN%\%USERINPUT%**.
- To append a domain to the user input, enter a **User Domain** and set the **Username Modifier** to **%USERINPUT%@%USERDOMAIN%**.

If the Username Modifier includes the %USERDOMAIN% variable, the User Domain value replaces any domain string that the user enters. If the User Domain is blank, that means the device removes any user-entered domain string.

8. Select the **Advanced** tab.
9. In the Allow List, **Add** and then select the users and groups that are allowed to authenticate with this profile. Selecting the predefined **all** option allows every user to authenticate. By default, the list has no entries, which means no users can authenticate.
10. Click **OK**.

STEP 3 | Commit the configuration.

Click **Commit**.

Set Up SAML Authentication

Security Assertion Markup Language (SAML) is an XML-based, open-standard data format for exchanging authentication and authorization data between parties, in particular, between an identity provider (IdP) and a service provider. SAML is a product of the OASIS Security Services Technical Committee.

STEP 1 | Create a server profile.

The server profile identifies the external authentication service and instructs the firewall how to connect to that authentication service and access the authentication credentials for your users.

1. Select **Device** > **Server Profiles** and select the **SAML Identity Provider** profile.
2. Click **Add** and enter a Profile **Name**, such as GP-User-Auth.
3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the Location where the profile is available.
4. **Import** the IdP metadata file. Refer to SAML 2.0 Authentication for details.

Alternatively, if the IdP doesn't provide a metadata file, Add the server profile and then enter the connection and registration information.

5. Click **OK** to save the server profile.

STEP 2 | (**Optional**) Create an authentication profile.

The authentication profile specifies the server profile for the portal or gateways to use when they authenticate users. On a portal or gateway, you can assign one or more authentication profiles in one or more *client authentication* profiles. For descriptions of how an authentication profile within a client authentication profile supports granular user authentication, see Configure a GlobalProtect Gateway and Set Up Access to the GlobalProtect Portal.

SAML authentication does not support the pre-logon connect method that enables users to connect and change their own expired passwords without administrative intervention (Remote Access VPN with Pre-Logon).

1. Select **Device** > **Authentication Profile** and **Add** a new profile.
2. Enter a **Name** for the profile and then select **SAML** as the authentication **Type**.
3. Select the SAML authentication **Server Profile** that you created in 1 from the drop-down.

© 2017 Palo Alto Networks, Inc.

4. Select the following to configure certificate authentication between the firewall and the SAML identity provider. Refer to SAML 2.0 Authentication for details.
   - The **Request Signing Certificate** that the firewall uses to sign messages it sends to the IdP.
   - The **Certificate Profile** that the firewall uses to validate the **Identity Provider Certificate**.
5. Specify the username and admin role formats.
   - Specify the **Username Attribute** and **User Group Attribute**.

   Unlike other types of external authentication, there is no User Domain attribute in the authentication profiles for SAML.

   - (**Optional**) If you will use this profile to authenticate administrative accounts that you manage in the IdP identity store, specify the **Admin Role Attribute** and **Access Domain Attribute** also.
6. Select the **Advanced** tab.
7. In the Allow List, **Add** and then select the users and groups that are allowed to authenticate with this profile. Selecting the predefined **all** option allows every user to authenticate. By default, the list has no entries, which means no users can authenticate.

   Make sure the username in the Allow List matches the username returned from the SAML IdP server.
8. Click **OK**.

STEP 3 | Commit the configuration.

Click **Commit**.

Set Up Kerberos Authentication

Kerberos is a computer network authentication protocol that works on the basis of *tickets* to allow nodes communicating over a non-secure network to prove their identity to one another in a secure manner.

STEP 1 | Create a server profile.

The server profile identifies the external authentication service and instructs the firewall how to connect to that authentication service and access the authentication credentials for your users.

1. Select **Device** > **Server Profiles** and select the **Kerberos** profile.
2. Click **Add** and enter a Profile **Name**, such as GP-User-Auth.
3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the Location where the profile is available.
4. Click **Add** in the Servers section and then enter the necessary information for connecting to the authentication server, including the server **Name**, IP address or FQDN of the **Server**, and **Port**.
5. Click **OK** to save the server profile.

STEP 2 | (**Optional**) Create an authentication profile.

The authentication profile specifies the server profile for the portal or gateways to use when they authenticate users. On a portal or gateway, you can assign one or more authentication profiles in one or more *client authentication* profiles. For descriptions of how an authentication profile within a client authentication profile supports granular user authentication, see Configure a GlobalProtect Gateway and Set Up Access to the GlobalProtect Portal.

To enable users to connect and change their own expired passwords without administrative intervention, consider using a pre-logon connect method. See Remote Access VPN with Pre-Logon *for details.*

© 2017 Palo Alto Networks, Inc.

1. Select **Device** > **Authentication Profile** and **Add** a new profile.
2. Enter a **Name** for the profile and then select **Kerberos** as the authentication **Type**.
3. Select the Kerberos authentication **Server Profile** that you created in 1 from the drop-down.
4. Specify the domain name and username format. The device combines the **User Domain** and **Username Modifier** values to modify the domain/username string that a user enters during login. The device uses the modified string for authentication and uses the **User Domain** value for User-ID group mapping. Modifying user input is useful when the authentication service requires domain/username strings in a particular format and you don't want to rely on users to correctly enter the domain. You can select from the following options:
   - To send only the unmodified user input, leave the **User Domain** blank (the default) and set the **Username Modifier** to the variable **%USERINPUT%** (the default).
   - To prepend a domain to the user input, enter a **User Domain** and set the **Username Modifier** to **%USERDOMAIN%\%USERINPUT%**.
   - To append a domain to the user input, enter a **User Domain** and set the **Username Modifier** to **%USERINPUT%@%USERDOMAIN%**.

   If the Username Modifier includes the %USERDOMAIN% variable, the User Domain value replaces any domain string that the user enters. If the User Domain is blank, that means the device removes any user-entered domain string.
5. Configure Kerberos single sign-on (SSO) if your network supports it:
   - Enter the **Kerberos Realm** (up to 127 characters). This is the hostname portion of the user login name. For example, the user account name user@EXAMPLE.LOCAL has the realm EXAMPLE.LOCAL.
   - Specify a **Kerberos Keytab** file: click the **Import** link, **Browse** to the keytab file, and click **OK**. During authentication, the endpoint first tries to use the keytab to establish SSO. If it succeeds, and the user attempting access is in the **Allow List**, authentication succeeds immediately. Otherwise, the authentication process falls back to manual (username/password) authentication of the specified **Type**. The **Type** doesn't have to be Kerberos. To change this behavior so that users can authenticate only using Kerberos, set **Use Default Authentication on Kerberos Authentication Failure** to **No** in a GlobalProtect portal agent configuration.
6. Select the **Advanced** tab.
7. In the Allow List, **Add** and then select the users and groups that are allowed to authenticate with this profile. Selecting the predefined **all** option allows every user to authenticate. By default, the list has no entries, which means no users can authenticate.

STEP 3 | Save your changes and commit the configuration.

Click **OK**, and then **Commit**.

Set Up RADIUS or TACACS+ Authentication

RADIUS is a client/server protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. TACACS+ is a well-established authentication protocol common to UNIX networks that allows a remote access server to forward a user's logon password to an authentication server to determine whether access can be allowed to a given system.

STEP 1 | Create a server profile.

The server profile identifies the external authentication service and instructs the firewall how to connect to that authentication service and access the authentication credentials for your users.

© 2017 Palo Alto Networks, Inc.

If you want to Enable Delivery of GlobalProtect Client VSAs to a RADIUS Server*, you must create a RADIUS server profile.*

1. Select **Device** > **Server Profiles** and select the type of profile (**RADIUS** or **TACACS+**).
2. Click **Add** and enter a Profile **Name**, such as GP-User-Auth.
3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the Location where the profile is available.
4. Configure the following Server Settings. These settings to all servers you include in the profile.
   - **Timeout (sec)**—The number of seconds before a server connection request times out due to lack of response from the authentication server.
   - **Authentication Protocol**—Select the protocol to use for connections to the authentication server. Choices are **CHAP**, **PAP**, or **Auto**.
   - (**RADIUS only**) **Retries**—The number of times the firewall trys connecting to the authentication server before dropping the request.
   - (**TACACS+ only**) **Use single connection for all authentication** to allow all TACACS+ authentication requests to occur over a single TCP session rather than separate sessions for each request.
5. Click **Add** in the Servers section and then enter the necessary information for connecting to the authentication server, including the server **Name**, IP address or FQDN of the **Server**, and **Port**.
6. Specify settings to enable the authentication service to authenticate the firewall. Enter the shared **Secret** when adding the server entry.
7. Click **OK** to save the server profile.

STEP 2 | (**Optional**) Create an authentication profile.

The authentication profile specifies the server profile for the portal or gateways to use when they authenticate users. On a portal or gateway, you can assign one or more authentication profiles in one or more *client authentication* profiles. For descriptions of how an authentication profile within a client authentication profile supports granular user authentication, see Configure a GlobalProtect Gateway and Set Up Access to the GlobalProtect Portal.

To enable users to connect and change their own expired passwords without administrative intervention, consider using a pre-logon connect method. See Remote Access VPN with Pre-Logon *for details.*

1. Select **Device** > **Authentication Profile** and **Add** a new profile.
2. Enter a **Name** for the profile and then select the authentication **Type** (**RADIUS** or **TACACS+**).
3. Select the RADIUS or TACACS+ authentication **Server Profile**. that you created in 1 from the drop-down.
4. (**RADIUS only**) Enable **Retrieve user group from RADIUS** if you want to include this information in the authentication profile.
5. Specify the domain name and username format. The device combines the **User Domain** and **Username Modifier** values to modify the domain/username string that a user enters during login. The device uses the modified string for authentication and uses the **User Domain** value for User-ID group mapping. Modifying user input is useful when the authentication service requires domain/username strings in a particular format and you don't want to rely on users to correctly enter the domain. You can select from the following options:
   - To send only the unmodified user input, leave the **User Domain** blank (the default) and set the **Username Modifier** to the variable **%USERINPUT%** (the default).
   - To prepend a domain to the user input, enter a **User Domain** and set the **Username Modifier** to **%USERDOMAIN%\%USERINPUT%**.

© 2017 Palo Alto Networks, Inc.

- To append a domain to the user input, enter a **User Domain** and set the **Username Modifier** to **%USERINPUT%@%USERDOMAIN%**.



If the Username Modifier includes the %USERDOMAIN% variable, the User Domain value replaces any domain string that the user enters. If the User Domain is blank, that means the device removes any user-entered domain string.

6. Select the **Advanced** tab.
7. In the Allow List, **Add** and then select the users and groups that are allowed to authenticate with this profile. Selecting the predefined **all** option allows every user to authenticate. By default, the list has no entries, which means no users can authenticate.
8. Click **OK**.

STEP 3 | Commit the configuration.

Click **Commit**.

© 2017 Palo Alto Networks, Inc.

Set Up Client Certificate Authentication

With the optional client certificate authentication, the agent/app presents a client certificate along with its connection request to the GlobalProtect portal or gateway. The portal or gateway can use either a shared or unique client certificate to validate that the user or device belongs to your organization.

The methods for deploying client certificates depend on the security requirements for your organization:

- Deploy Shared Client Certificates for Authentication on page 38
- Deploy Machine Certificates for Authentication on page 38
- Deploy User-Specific Client Certificates for Authentication on page 42

Deploy Shared Client Certificates for Authentication

To confirm that an endpoint user belongs to your organization, you can use the same client certificate for all endpoints or generate separate certificates to deploy with a particular agent configuration. Use this workflow to issue self-signed client certificates for this purpose and deploy them from the portal.

STEP 1 | Generate a certificate to deploy to multiple GlobalProtect clients.

1. Create the root CA certificate for issuing self-signed certificates for the GlobalProtect components.
2. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**.
3. Use the **Local** certificate type (the default).
4. Enter a **Certificate Name**. This name cannot contain spaces.
5. In the **Common Name** field enter a name to identify this certificate as an agent certificate, for example GP_Windows_clients. Because this same certificate will be deployed to all agents using the same configuration, it does not need to uniquely identify a specific user or endpoint.
6. In the **Signed By** field, select your root CA.
7. Select an **OCSP Responder** to verify the revocation status of certificates.
8. Click **OK** to generate the certificate.

STEP 2 | Set Up Two-Factor Authentication.

Configure authentication settings in a GlobalProtect portal agent configuration to enable the portal to transparently deploy the client certificate that is **Local** to the firewall to clients that receive the configuration.

Deploy Machine Certificates for Authentication

To confirm that the endpoint belongs to your organization, use your own public-key infrastructure (PKI) to issue and distribute machine certificates to each endpoint (recommended) or generate a self-signed machine certificate for export. With the pre-logon connect methods, a machine certificate is required and must be installed on the endpoint before GlobalProtect components will grant access.

To confirm that the endpoint belongs to your organization, you must also configure an authentication profile to authenticate the user. See Two-Factor Authentication.

Use the following workflow to create the client certificate and manually deploy it to an endpoint. For more information, see About GlobalProtect User Authentication. For an example configuration, see Remote Access VPN (Certificate Profile).

STEP 1 | Issue client certificates to GlobalProtect clients and endpoints.

© 2017 Palo Alto Networks, Inc.

This enables the GlobalProtect portal and gateways to validate that the device belongs to your organization.

1. Create the root CA certificate for issuing self-signed certificates for the GlobalProtect components.
2. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**.
3. Enter a **Certificate Name**. The certificate name cannot contain any spaces.
4. Configure cryptographic settings for the certificate including the encryption **Algorithm**, key length (**Number of Bits**), **Digest** algorithm (use sha1, sha256, or sha384; sha512 is not supported with client certificates), and **Expiration** (in days) for the certificate.

   If the firewall is in FIPS-CC mode and the key generation algorithm is RSA, the RSA keys must be 2,048 bits or 3072 bits.
5. In the Certificate Attributes section, **Add** and define the attributes that uniquely identify the GlobalProtect clients as belonging to your organization. Keep in mind that if you add a **Host Name** attribute (which populates the SAN field of the certificate), it must be the same as the value you defined for the **Common Name**.
6. In the **Signed By** field, select your root CA.
7. Select an **OCSP Responder** to verify the revocation status of certificates.
8. (**Optional**) In the Certificate Attributes section, click **Add** and define the attributes to identify the GlobalProtect clients as belonging to your organization if required as part of your security requirements.
9. Click **OK** to generate the certificate.

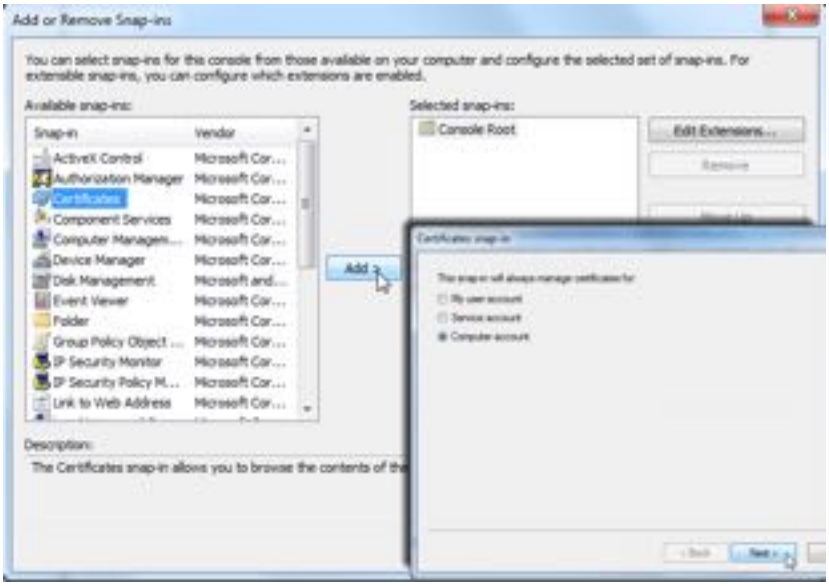
STEP 2 | Install certificates in the personal certificate store on the endpoints.

If you are using unique user certificates or machine certificates, you must install each certificate in the personal certificate store on the endpoint prior to the first portal or gateway connection. Install machine certificates to the Local Computer certificate store on Windows and in the System Keychain on Mac OS. Install user certificates to the Current User certificate store on Windows and in the Keychain on Mac OS.

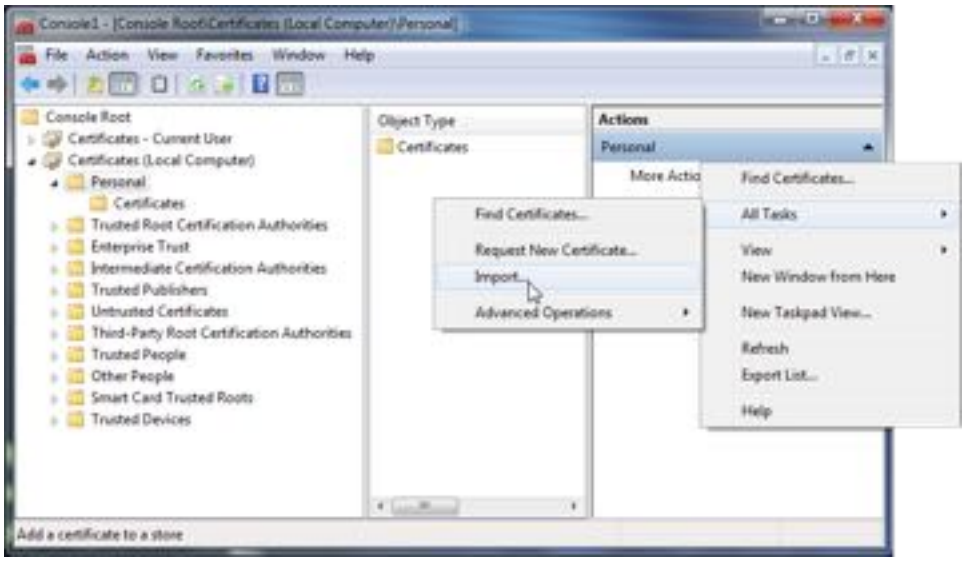
For example, to install a certificate on a Windows system using the Microsoft Management Console:

1. From the command prompt, enter `mmc` to launch the console.
2. Select **File** > **Add/Remove Snap-in**.
3. Select **Certificates**, click **Add** and then select one of the following, depending on what type of certificate you are importing:
   - **Computer account**—Select this option if you are importing a machine certificate.
   - **My user account**—Select this option if you are importing a user certificate.

© 2017 Palo Alto Networks, Inc.



4. Expand **Certificates** and select **Personal** and then in the Actions column select **Personal** > **More Actions** > **All Tasks** > **Import** and follow the steps in the Certificate Import Wizard to import the PKCS file you got from the CA.

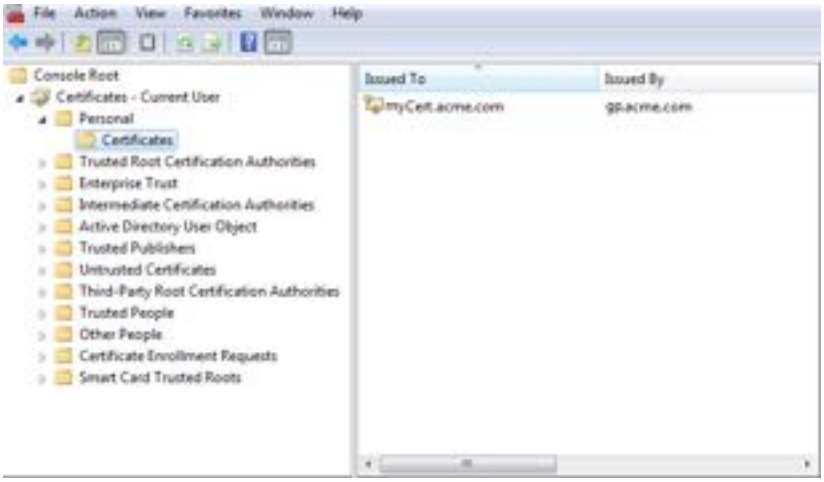


5. Browse to the .p12 certificate file to import (select **Personal Information Exchange** as the file type to browse for) and enter the **Password** that you used to encrypt the private key. Select **Personal** as the **Certificate store**.

STEP 3 | Verify that the certificate has been added to the personal certificate store.

Navigate to the personal certificate store:

© 2017 Palo Alto Networks, Inc.



STEP 4 | Import the root CA certificate used to issue the client certificates onto the firewall.

This step is required only if an external CA issued the client certificates, such as a public CA or an enterprise PKI CA. If you are using self-signed certificates, the root CA is already trusted by the portal and gateways.

1. Download the root CA certificate used to issue the client certificates (Base64 format).
2. Import the root CA certificate from the CA that generated the client certificates onto the firewall:
    1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
    2. Use the **Local** certificate type (the default).
    3. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.
    4. **Browse** to the **Certificate File** you downloaded from the CA.
    5. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.
    6. Select the certificate you just imported on the **Device Certificates** tab to open it.
    7. Select **Trusted Root CA** and then click **OK**.

STEP 5 | Create a client certificate profile.

1. Select **Device** > **Certificates** > **Certificate Management** > **Certificate Profile**, click **Add**, and enter a profile **Name**.
2. Select a value for the **Username Field** to specify which field in the certificate will contain the user's identity information.

    If you plan to configure the portal or gateways to authenticate users with certificates only, you must specify the **Username Field**. This enables GlobalProtect to associate a username with the certificate.

    If you plan to set up the portal or gateway for two-factor authentication, you can leave the default value of **None**, or, to add an additional layer of security, specify a username. If you specify a username, your external authentication service verifies that the username in the client certificate matches the username requesting authentication. This ensures that the user is the one to which the certificate was issued.

    Users cannot change the username that is included in the certificate.
3. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you imported in 4 and then click **OK**.

STEP 6 | Save the configuration.

Click **Commit**.

© 2017 Palo Alto Networks, Inc.

Deploy User-Specific Client Certificates for Authentication

To authenticate individual users, you must issue a unique client certificate to each GlobalProtect user and deploy the client certificate to the endpoints prior to enabling GlobalProtect. To automate the generation and deployment of user-specific client certificates, you can configure your GlobalProtect portal to act as a Simple Certificate Enrollment Protocol (SCEP) client to a SCEP server in your enterprise PKI.

SCEP operation is dynamic in that the enterprise PKI generates a user-specific certificate when the portal requests it and sends the certificate to the portal. The portal then transparently deploys the certificate to the client. When a user requests access, the agent or app can then present the client certificate to authenticate with the portal or gateway.

The GlobalProtect portal or gateway uses identifying information about the device and user to evaluate whether to permit access to the user. GlobalProtect blocks access if the host ID is on a device block list or if the session matches any blocking options specified in a certificate profile. If client authentication fails due to an invalid SCEP-based client certificate, the GlobalProtect client tries to authenticate with the portal per the settings in the authentication profile and retrieve the certificate. If the client cannot retrieve the certificate from the portal, the device is not able to connect.

STEP 1 | Create a SCEP profile.

1. Select **Device** > **Certificate Management** > **SCEP** and then **Add** a new profile.
2. Enter a **Name** to identify the SCEP profile.
3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the **Location** where the profile is available.

STEP 2 | (**Optional**) To make the SCEP-based certificate generation more secure, configure a SCEP challenge-response mechanism between the PKI and portal for each certificate request.

After you configure this mechanism, its operation is invisible, and no further input from you is necessary.

To comply with the U.S. Federal Information Processing Standard (FIPS), use a **Dynamic** SCEP challenge and specify a **Server URL** that uses HTTPS (see 7).

Select one of the following options:

- **None**—(**Default**) The SCEP server does not challenge the portal before it issues a certificate.
- **Fixed**—Obtain the enrollment challenge password from the SCEP server in the PKI infrastructure and then enter the password into the Password field.
- **Dynamic**—Enter a username and password of your choice (possibly the credentials of the PKI administrator) and the SCEP **Server URL** where the portal-client submits these credentials. The uses the credentials to authenticate with the SCEP server which transparently generates an OTP password for the portal upon each certificate request. (You can see this OTP change after a screen refresh in `The enrollment challenge password is` field after each certificate request.) The PKI transparently passes each new password to the portal, which then uses the password for its certificate request.

STEP 3 | Specify the settings for the connection between the SCEP server and the portal to enable the portal to request and receive client certificates.

You can include additional information about the client device or user by specifying tokens in the **Subject** name of the certificate.

In the Subject field of the CSR to the SCEP server, the portal includes the token value as CN and Host-ID as SerialNumber. The host ID varies by device type: GUID (Windows) MAC address of the interface (Mac), Android ID (Android devices), UDID (iOS devices), or a unique name that GlobalProtect assigns (Chrome).

© 2017 Palo Alto Networks, Inc.

1. Configure the **Server URL** that the portal uses to reach the SCEP server in the PKI (for example, `http://10.200.101.1/certsrv/mscep/`).
2. Enter a string (up to 255 characters in length) in the **CA-IDENT Name** field to identify the SCEP server.
3. Enter the **Subject** name to use in the certificate generated by the SCEP server. The subject must be a distinguished name in the ***<attribute>=<value>*** format and must include a common name (CN) attribute (**CN=<*token*>**). The CN supports the following dynamic tokens:
   - **$USERNAME**—Use this token to enable the portal to request certificates for a specific user. To use this variable, you must also Enable Group Mapping. The username entered by the user must match the name in the user-group mapping table.
   - **$EMAILADDRESS**—Use this token to request certificates associated with a specific email address. To use this variable, you must also Enable Group Mapping and configure the **Mail Attributes** in the Mail Domains section of the Server Profile. If GlobalProtect cannot identify an email address for the user, it generates a unique ID and populates the CN with that value.
   - **$HOSTID**—To request certificates for the device only, specify the host ID token. When a user attempts to log in to the portal, the endpoint sends identifying information that includes its host ID value.

   When the GlobalProtect portal pushes the SCEP settings to the agent, the CN portion of the subject name is replaced with the actual value (username, host ID, or email address) of the certificate owner (for example, **O=acme,CN=johndoe**).
4. Select the **Subject Alternative Name Type**:
   - **RFC 822 Name**—Enter the email name in a certificate's subject or Subject Alternative Name extension.
   - **DNS Name**—Enter the DNS name used to evaluate certificates.
   - **Uniform Resource Identifier**—Enter the name of the resource from which the client will obtain the certificate.
   - **None**—Do not specify attributes for the certificate.

STEP 4 | (**Optional**) Configure cryptographic settings for the certificate.

- Select the key length (**Number of Bits**) for the certificate.

  If the firewall is in FIPS-CC mode and the key generation algorithm is RSA. The RSA keys must be 2,048 bits or larger.
- Select the **Digest for CSR** which indicates the digest algorithm for the certificate signing request (CSR): sha1, sha256, or sha384.



Sha512 is not supported as a digest algorithm for client certificates on GlobalProtect endpoints.

STEP 5 | (**Optional**) Configure the permitted uses of the certificate, either for signing or encryption.

- To use this certificate for signing, select the **Use as digital signature** check box. This enables the endpoint use the private key in the certificate to validate a digital signature.
- To use this certificate for encryption, select the **Use for key encipherment** check box. This enables the client use the private key in the certificate to encrypt data exchanged over the HTTPS connection established with the certificates issued by the SCEP server.

STEP 6 | (**Optional**) To ensure that the portal is connecting to the correct SCEP server, enter the **CA Certificate Fingerprint**. Obtain this fingerprint from the SCEP server interface in the Thumbprint field.

© 2017 Palo Alto Networks, Inc.

1. Enter the URL for the SCEP server's administrative UI (for example, **`http://<hostname or IP>/CertSrv/mscep_admin/`**).
2. Copy the thumbprint and enter it in the **CA Certificate Fingerprint** field.

STEP 7 | Enable mutual SSL authentication between the SCEP server and the GlobalProtect portal. This is required to comply with the U.S. Federal Information Processing Standard (FIPS).



FIPS-CC operation is indicated on the firewall login page and in its status bar.

Select the SCEP server's root **CA Certificate**. Optionally, you can enable mutual SSL authentication between the SCEP server and the GlobalProtect portal by selecting a **Client Certificate**.

STEP 8 | Save and commit the configuration.

1. Click **OK** to save the settings and close the SCEP configuration.
2. **Commit** the configuration.

The portal attempts to request a CA certificate using the settings in the SCEP profile and saves it to the firewall hosting the portal. If successful, the CA certificate is shown in **Device** > **Certificate Management** > **Certificates**.

STEP 9 | (**Optional**) If after saving the SCEP profile, the portal fails to obtain the certificate, you can manually generate a certificate signing request (CSR) from the portal.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**.
2. Enter a **Certificate Name**. This name cannot contain spaces.
3. Select the **SCEP Profile** to use to submit a CSR to your enterprise PKI.
4. Click **OK** to submit the request and generate the certificate.

STEP 10 | Set Up Two-Factor Authentication.

Assign the SCEP profile a GlobalProtect portal agent configuration to enable the portal to transparently request and deploy client certificates to clients that receive the configuration.

© 2017 Palo Alto Networks, Inc.

Set Up Two-Factor Authentication

If you require strong authentication to protect sensitive assets or to comply with regulatory requirements, such as PCI, SOX, or HIPAA, configure GlobalProtect to use an authentication service that uses a two-factor authentication scheme. A two-factor authentication scheme requires two things: something the end user knows (such as a PIN or password) and something the end user has (a hardware or software token/OTP, smart card, or certificate). You can also enable two-factor authentication using a combination of external authentication services, and client and certificate profiles.

The following topics provide examples for how to set up two-factor authentication on GlobalProtect:

- Enable Two-Factor Authentication Using Certificate and Authentication Profiles on page 45
- Enable Two-Factor Authentication Using One-Time Passwords (OTPs) on page 47
- Enable Two-Factor Authentication Using Smart Cards on page 50

Enable Two-Factor Authentication Using Certificate and Authentication Profiles

The following workflow describes how to configure GlobalProtect client authentication requiring the user to authenticate both to a certificate profile and an authentication profile. The user must successfully authenticate using both methods in order to connect to the portal/gateway. For more details on this configuration, see Remote Access VPN with Two-Factor Authentication.

STEP 1 | Create an authentication server profile.

The authentication server profile determines how the firewall connects to an external authentication service and retrieves the authentication credentials for your users.



If you are using LDAP to connect to Active Directory (AD), you must create a separate LDAP server profile for every AD domain.

1. Select **Device** > **Server Profiles** and a profile type (**LDAP**, **Kerberos**, **RADIUS**, or **TACACS+**).
2. **Add** a new server profile.
3. Enter a **Profile Name** for the profile, such as GP-User-Auth.
4. (**LDAP only**) Select the **Type** of LDAP server (**active-directory**, **e-directory**, **sun**, or **other**).
5. Click **Add** in the Servers list section and then enter the required information for connections to the authentication service, including the server **Name**, IP address or FQDN of the **Server**, and **Port**.
6. (**RADIUS, TACACS+, and LDAP only**) Specify settings to enable the firewall to authenticate to the authentication service as follows:
   - RADIUS and TACACS+—Enter the shared **Secret** when adding the server entry.
   - LDAP—Enter the **Bind DN** and **Password**.
7. (**LDAP only**) If you want the endpoint to use SSL or TLS for a more secure connection with the directory server, select the **Require SSL/TLS secured connection** check box (selected by default). The protocol that the device uses depends on the server **Port** in the **Server list**:
   - 389 (default)—TLS (specifically, the device uses the StartTLS operation to upgrade the initial plaintext connection to TLS).
   - 636—SSL.
   - Any other port—The device first attempts to use TLS. If the directory server does not support TLS, the device uses SSL.
8. (**LDAP only**) For additional security, select the **Verify Server Certificate for SSL sessions** check box so that the endpoint verifies the certificate that the directory server presents for SSL/TLS connections.

© 2017 Palo Alto Networks, Inc.

To enable verification, you also must select the **Require SSL/TLS secured connection** check box. For verification to succeed, one of the following conditions must be true:

- The certificate is in the list of device certificates: **Device** > **Certificate Management** > **Certificates** > **Device Certificates**. Import the certificate into the endpoint if necessary.
- The certificate signer is in the list of trusted certificate authorities: **Device** > **Certificate Management** > **Certificates** > **Default Trusted Certificate Authorities**.

9. Click **OK** to save the server profile.

STEP 2 | Create an authentication profile that identifies the service for authenticating users. (You later have the option of assigning the profile on the portal and on gateways.)

1. Select **Device** > **Authentication Profile** and **Add** a new profile.
2. Enter a **Name** for the profile.
3. Select the **Location**.
4. Select the **Type** of **Authentication** (**LDAP**, **Kerberos**, **RADIUS**, or **TACACS+**).
5. Select the **Server Profile** you created in 1.
6. (**LDAP only**) Enter **sAMAccountName** as the **Login Attribute**.
7. Click **OK** to save the authentication profile.

STEP 3 | Create a client certificate profile that the portal uses to authenticate the client certificates that come from user devices.



When you configure two-factor authentication to use client certificates, the external authentication service uses the username value to authenticate the user, if specified, in the client certificate. This ensures that the user who is logging is in is actually the user to whom the certificate was issued.

1. Select **Device** > **Certificates** > **Certificate Management** > **Certificate Profile** and click **Add** and enter a profile **Name**.
2. Select a value for the **Username Field**:
   - If you intend for the client certificate to authenticate individual users, select the certificate field that identifies the user.
   - If you are deploying the client certificate from the portal, leave this field set to **None**.
   - If you are setting up a certificate profile for use with a pre-logon connect method, leave the field set to **None**.
3. In the **CA Certificates** area, click **Add** and then:
   1. Select the **CA certificate**, either a trusted root CA certificate or the CA certificate from a SCEP server. (If necessary, import the certificate).
   2. (**Optional**) Enter the **Default OCSP URL**.
   3. (**Optional**) Select a certificate for **OCSP Verify CA**.
4. (**Optional**) Select options that specify when to block the user's requested session:
   1. Status of certificate is unknown.
   2. GlobalProtect component does not retrieve certificate status within the number of seconds in **Certificate Status Timeout**.
   3. Serial number attribute in the subject of a client certificate does not match the host ID that the GlobalProtect agent reports for the client endpoint.
5. Click **OK**.

STEP 4 | (**Optional**) Issue client certificates to GlobalProtect clients and endpoints.

© 2017 Palo Alto Networks, Inc.

To transparently deploy client certificates, configure your portal to distribute a shared client certificate to your endpoints or configure the portal to use SCEP to request and deploy unique client certificates for each user.

1. Use your enterprise PKI or a public CA to issue a client certificate to each GlobalProtect user.
2. For the pre-logon connect methods, install certificates in the personal certificate store on the client systems.

STEP 5 | Save the GlobalProtect configuration.

Click **Commit**.

Enable Two-Factor Authentication Using One-Time Passwords (OTPs)

Use this workflow to configure two-factor authentication using one-time passwords (OTPs) on the portal and gateways. When a user requests access, the portal or gateway prompts the user to enter an OTP. The authentication service sends the OTP as a token to the user's RSA device.

Setting up a two-factor authentication scheme is similar to setting up other types of authentication and requires you to configure:

- A server profile (usually for a RADIUS service for two-factor authentication) assigned to an authentication profile.
- A client authentication profile that includes the authentication profile for the service that these components use.

By default, the agent supplies the same credentials it used to log in to the portal and to the gateway. In the case of OTP authentication, this behavior will cause the authentication to initially fail on the gateway and, because of the delay this causes in prompting the user for a login, the user's OTP may expire. To prevent this, you must configure the portals and gateways that prompt for the OTP instead of using the same credentials on a per-agent configuration basis.

You can also reduce the frequency in which users are prompted for OTPs by configuring an authentication override. This enables the portals and gateways to generate and accept a secure encrypted cookie to authenticate the user for a specified amount of time. The portals and/or gateways will not require a new OTP until the cookie expires thus reducing the number of times users must provide an OTP.

STEP 1 | After you have configured the back-end RADIUS service to generate tokens for the OTPs and ensured users have any necessary devices (such as a hardware token), set up a RADIUS server to interact with the firewall.

For specific instructions, refer to the documentation for your RADIUS server. In most cases, you need to set up an authentication agent and a client configuration on the RADIUS server to enable communication between the firewall and the RADIUS server. You also define the shared secret to use for encrypting sessions between the firewall and the RADIUS server.

STEP 2 | On each firewall that hosts the gateways and/or portal, create a RADIUS server profile. (For a small deployment, one firewall can host the portal and gateways.)

1. Select **Device** > **Server Profiles** > **RADIUS**.
2. **Add** a new profile.
3. Enter a **Name** for this RADIUS profile.
4. In the **Servers** area, **Add** a RADIUS instance and enter:
    - A descriptive **Name** to identify this RADIUS server

© 2017 Palo Alto Networks, Inc.

- The **RADIUS Server** IP address
- The shared **Secret** for encrypting sessions between the firewall and the RADIUS server
- The **Port** number on which the RADIUS server listens for authentication requests (default 1812)

5. Click **OK** to save the profile.

STEP 3 | Create an authentication profile.

1. Select **Device** > **Authentication Profile**.
2. **Add** a new profile.
3. Enter a **Name** for the profile. The name cannot contain spaces.
4. Select **RADIUS** as the **Type** of authentication service.
5. Select the **Server Profile** you created for accessing your RADIUS server.
6. Enter the **User Domain** name. The firewall uses this value for matching authenticating users against Allow List entries and for User-ID group mapping.
7. Select a **Username Modifier** to modify the username/domain format expected by the RADIUS server.
8. Click **OK** to save the authentication profile.

STEP 4 | Assign the authentication profile to the GlobalProtect gateway(s) and/or portal.

You can configure multiple Client Authentication configurations for the portal and gateways. For each Client Authentication configuration you can specify the authentication profile to apply to endpoints of a specific OS.

This step describes only how to add the authentication profile to the gateway or portal configuration. For additional details on setting up these components, see GlobalProtect Gateways and GlobalProtect Portals.

1. Select **Network** > **GlobalProtect** > **Gateways** and an existing gateway configuration by name (or **Add** one). If you are adding a new gateway, specify its name, location, and network parameters.
2. On the **Authentication** tab, select an SSL/TLS service profile or **Add** a new profile.
3. **Add** a Client Authentication configuration and enter its **Name**.
4. Select the endpoint OS to which this configuration applies.
5. Select the **Authentication Profile** you created in Create an authentication profile.
6. (**Optional**) Enter a custom authentication message.
7. To add additional Client Authentication configurations, repeat steps c through f.
8. Click **OK** to save the configuration.
9. To add other gateways, repeat steps b through h.
10. To assign the authentication profile to the portal, select **Network** > **GlobalProtect** > **Portals** and repeat steps b through h.

STEP 5 | (**Optional**) Configure the portal or gateways to prompt for a username and password or only a password each time the user logs in. Saving the password is not supported with two-factor authentication using OTPs because the user must enter a dynamic password each time they log in.

This step describes only how to configure the password setting in a portal agent configuration. For additional details, see Customize the GlobalProtect Agent.

1. Select **Network** > **GlobalProtect** > **Portals** and select an existing portal configuration.
2. Select **Agent**.
3. Select an existing agent configuration or **Add** one.
4. Set **Save User Credentials** to **Save Username Only** or **No**. This setting enables GlobalProtect to prompt for dynamic passwords for each component you select in the following step.
5. Click **OK** twice to save the configuration.

© 2017 Palo Alto Networks, Inc.

STEP 6 | Select the GlobalProtect components—portal and types of gateways—that prompt for dynamic passwords, such as OTPs, instead of using saved credentials.

1. Select **Network** > **GlobalProtect** > **Portals** and select an existing portal configuration.
2. Select **Agent**.
3. Select an existing agent configuration or **Add** one.
4. Select the **Authentication** tab, and then select the Components that Require Dynamic Passwords (Two-Factor Authentication). When selected, the portal and/or types of gateways prompt for OTPs.
5. Click **OK** twice to save the configuration.

STEP 7 | If single sign-on (SSO) is enabled, disable it. The agent configuration specifies RADIUS as the authentication service so Kerberos SSO is not supported.

This step describes only how to disable SSO. For more details, see Define the GlobalProtect Agent Configurations.

1. Select **Network** > **GlobalProtect** > **Portals** and select the portal configuration.
2. Select **Agent** and then select the agent configuration (or **Add** one).
3. Select the **App** tab.
4. Set **Use Single Sign-on** to **No**.
5. Click **OK** twice to save the configuration.

STEP 8 | (**Optional**) To minimize the number of times a user must provide credentials, configure an authentication override.

By default, the portal or gateways authenticate the user with an authentication profile and optional certificate profile. With authentication override, the portal or gateway authenticates the user with an encrypted cookie that it has deployed to the endpoint. While the cookie is valid, the user can log in without entering regular credentials or an OTP. For more information, see Cookie Authentication on the Portal or Gateway.



If you need to immediately block access to a device whose cookie has not yet expired (for example, if the device is lost or stolen), you can Block Device Access by adding the device to a block list.

For more details, see GlobalProtect Gateways and GlobalProtect Portals.

1. Select **Network** > **GlobalProtect** > **Gateways** or **Portals** and select the configuration (or **Add** one).
2. Select **Agent** > **Client Settings** (on the gateway) or **Agent** (on the portal) and then select the configuration (or **Add** one).
3. In the **Authentication Override** area, configure the following:
   - **Generate cookie for authentication override**—Enable the portal or gateway to generate encrypted, endpoint-specific cookies. After users successfully authenticate, the portal or gateway issue the authentication cookie to the endpoint.
   - **Accept cookie for authentication override**—Select the check box to instruct the portal or gateway to authenticate the user through a valid, encrypted cookie. When the endpoint presents a valid cookie, the portal or gateway verifies that the cookie was encrypted by the portal or gateway, decrypts the cookie, and then authenticates the user.
   - **Cookie Lifetime**—Specify the hours, days, or weeks that the cookie is valid. Typical lifetime is 24 hours for gateways—which protect sensitive information—or 15 days for the portal. The range for hours is 1–72; for weeks, 1–52; and for days, 1–365. After the cookie expires on either the portal or gateway (whichever occurs first), the portal or gateway prompts the user to authenticate and subsequently encrypts a new cookie to send to the endpoint.
   - **Certificate to Encrypt/Decrypt Cookie**—Select the RSA certificate to use to encrypt and decrypt the cookie. You must use the same certificate on the portal and gateways.

© 2017 Palo Alto Networks, Inc.

As a best practice, configure the RSA certificate to use the strongest digest algorithm that your network supports.

The portal and gateways use the RSA encrypt padding scheme PKCS#1 V1.5 to generate the cookie (using the public key of the certificate) and decrypt the cookie (using the private key of the certificate).

4. Click **OK** twice to save the configuration.

STEP 9 | Commit the configuration.

Click **Commit**.

STEP 10 | Verify the configuration.

The gateway and portal must be configured before you take his step. For details on setting up these components, see GlobalProtect Gateways and GlobalProtect Portals.

From an endpoint running the GlobalProtect agent, try to connect to a gateway or portal on which you enabled OTP authentication. You should see two prompts similar to the following:

The first prompt requests a PIN (either a user- or system-generated PIN):



The second prompt requests your token or OTP:



Enable Two-Factor Authentication Using Smart Cards

If you want to enable your end users to authenticate using a smart card or common access card (CAC), you must import the Root CA certificate that issued the certificates contained on the end user CAC or smart

© 2017 Palo Alto Networks, Inc.

cards onto the portal and gateway. You can then create a certificate profile that includes that Root CA and apply it to your portal and/or gateway configurations to enable use of the smart card in the authentication process.

STEP 1 | Set up your smart card infrastructure.

This procedure assumes that you have deployed smart cards and smart card readers to your end users.

For specific instructions, refer to the documentation for the user authentication provider software.

In most cases, setting up the smart card infrastructure involves the generating of certificates for end users and for the participating servers, which are the GlobalProtect portal and gateway(s) in this use case.

STEP 2 | Import the Root CA certificate that issued the client certificates contained on the end user smart cards.

Make sure the certificate is accessible from your management system and then complete the following steps:

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates.**
2. Click **Import** and enter a **Certificate Name**.
3. Enter the path and name to the **Certificate File** received from the CA, or **Browse** to find the file.
4. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK** to import the certificate.

STEP 3 | Create the certificate profile on each portal/gateway on which you plan to use CAC or smart card authentication.



For details on other certificate profile fields, such as whether to use CRL or OCSP, refer to the online help.

1. Select **Device** > **Certificate Management** > **Certificate Profile** and click **Add** and enter a profile **Name**.
2. In the **Username** field, select the certificate field that PAN-OS uses to match the IP address for User-ID, either **Subject** to use a common name, **Subject Alt: Email** to use an email address, or **Subject Alt: Principal Name** to use the Principal Name.
3. In the **CA Certificates** field, click **Add**, select the trusted root **CA Certificate** you imported in 2 and then click **OK**.
4. Click **OK** to save the certificate profile.

STEP 4 | Assign the certificate profile to the gateway(s) or portal. This section describes only how to add the certificate profile to the gateway or portal configuration. For details on setting up these components, see GlobalProtect Gateways and GlobalProtect Portals.

1. Select **Network** > **GlobalProtect** > **Gateways** or **Portals** and select the configuration (or **Add** a new one).
2. On the **Authentication** tab, select the **Certificate Profile** you just created.
3. Click **OK** to save the configuration.

STEP 5 | Save the configuration.

Click **Commit**.

STEP 6 | Verify the configuration.

© 2017 Palo Alto Networks, Inc.

From a client system running the GlobalProtect agent, try to connect to a gateway or portal on which you set up smart card-enabled authentication. When prompted, insert your smart card and verify that you can successfully authenticate to GlobalProtect.

© 2017 Palo Alto Networks, Inc.

Set Up Authentication for strongSwan Ubuntu and CentOS Clients

To extend GlobalProtect VPN remote access support to strongSwan Ubuntu and CentOS clients, set up authentication for the strongSwan clients.

To view the minimum GlobalProtect release version that supports strongSwan on Ubuntu Linux and CentOS, see What Client OS Versions are Supported with GlobalProtect? *on page 11.*

To connect to the GlobalProtect gateway, the user must successfully authenticate. The following workflows show examples of how to enable authentication for strongSwan clients. For complete information about strongSwan, see the strongSwan wiki.

- Enable Authentication Using a Certificate Profile on page 53
- Enable Authentication Using an Authentication Profile on page 55
- Enable Authentication Using Two-Factor Authentication on page 57

Enable Authentication Using a Certificate Profile

The following workflow shows how to enable authentication for strongSwan clients using a certificate profile.

STEP 1 | Configure an IPSec tunnel for the GlobalProtect gateway for communicating with a strongSwan client.

1. Select **Network** > **GlobalProtect** > **Gateways** and then select the gateway name.
2. Select the **Certificate Profile** you want to use for authentication in the **Authentication** tab.
3. Select **Agent** > **Tunnel Settings** and specify the following settings to set up a tunnel:
   - Select the check box to **Enable X-Auth Support**.
   - If a **Group Name** and **Group Password** are already configured, remove them.
   - Click **OK** to save the settings.

STEP 2 | Verify that the default connection settings in the **`conn %default`** section of the IPSec tunnel configuration file (**`ipsec.conf`**) are correctly defined for the strongSwan client.

The `ipsec.conf` file is usually found in the `/etc` folder.

The configurations in this procedure are tested and verified for the following releases:

- *Ubuntu 14.0.4 with strongSwan 5.1.2 and CentOS 6.5 with strongSwan 5.1.3 for PAN-OS 6.1.*
- *Ubuntu 14.0.4 with strongSwan 5.2.1 for PAN-OS 7.0.*

The configurations in this procedure can be used for reference if you are using a different version of strongSwan. Refer to the strongSwan wiki for more information.

Modify the following settings in the `conn %default` section of the `ipsec.conf` file to these recommended settings.

```
ikelifetime=20m
reauth=yes
```


© 2017 Palo Alto Networks, Inc.

```
rekey=yes
keylife=10m
rekeymargin=3m
rekeyfuzz=0%
keyingtries=1
type=tunnel
```

STEP 3 | Modify the strongSwan client's IPSec configuration file (**`ipsec.conf`**) and the IPSec password file (**`ipsec.secrets`**) to use recommended settings.

The `ipsec.secrets` file is usually found in the `/etc` folder.

Use the strongSwan client username as the certificate's common name.

Modify the following items in the `ipsec.conf` file to these recommended settings.

```
conn <connection name>
keyexchange=ikev1
authby=rsasig
ike=aes-sha1-modp1024,aes256
left=<strongSwan/Linux-client-IP-address>
leftcert=<client certificate with the strongSwan client username used as the
 certificate's common name>
leftsourceip=%config
leftauth2=xauth
right=<GlobalProtect-Gateway-IP-address>
rightid="CN=<Subject-name-of-gateway-certificate>"
rightsubnet=0.0.0.0/0
auto=add
```

Modify the following items in the `ipsec.conf` file to these recommended settings.

```
:RSA
<private key file> "<passphrase if used>"
```

STEP 4 | Start strongSwan IPSec services and connect to the IPSec tunnel that you want the strongSwan client to use when authenticating to the GlobalProtect gateway.

Use the `config <name>` variable to name the tunnel configuration.

- Ubuntu clients:

```
ipsec start
ipsec up <name>
```

- CentOS clients:

```
strongSwan start
strongswan up <name>
```

STEP 5 | Verify that the tunnel is set up correctly and the VPN connection is established to both the strongSwan client and the GlobalProtect gateway.

1. Verify the detailed status information on a specific connection (by naming the connection) or verify the status information for all connections from the strongSwan client:
   - Ubuntu clients:

© 2017 Palo Alto Networks, Inc.

```
ipsec statusall [<connection name>]
```

- CentOS clients:

```
strongswan statusall [<connection name>]
```

2. Select **Network** > **GlobalProtect** > **Gateways**. Then, in the Info column, select **Remote Users** for the gateway configured for the connection to the strongSwan client. The strongSwan client should be listed under **Current Users**.

Enable Authentication Using an Authentication Profile

The following workflow shows how to enable authentication for strongSwan clients using an authentication profile. The authentication profile specifies which server profile to use when authenticating strongSwan clients.

STEP 1 | Set up the IPSec tunnel that the GlobalProtect gateway will use for communicating with a strongSwan client.

1. Select **Network** > **GlobalProtect** > **Gateways** and select the gateway name.
2. Select the **Authentication Profile** you want to use in the **Authentication** tab.
3. Select **Agent** > **Tunnel Settings** and specify the following settings to set up a tunnel:
   - Select the check box to **Enable X-Auth Support**.
   - Enter a **Group Name** and **Group Password** if they are not already configured.
   - Click **OK** to save these tunnel settings.

STEP 2 | Verify that the default connection settings in the **`conn %default`** section of the IPSec tunnel configuration file (**`ipsec.conf`**) are correctly defined for the strongSwan client.

The `ipsec.conf` file is usually found in the `/etc` folder.



The configurations in this procedure are tested and verified for the following releases:

- *Ubuntu 14.0.4 with strongSwan 5.1.2 and CentOS 6.5 with strongSwan 5.1.3 for PAN-OS 6.1.*
- *Ubuntu 14.0.4 with strongSwan 5.2.1 for PAN-OS 7.0.*

The configurations in this procedure can be used for reference if you are using a different version of strongSwan. Refer to the strongSwan wiki *for more information.*

In the `conn %default` section of the `ipsec.conf` file, configure the following recommended settings:

```
ikelifetime=20m
reauth=yes
rekey=yes
keylife=10m
rekeymargin=3m
rekeyfuzz=0%
keyingtries=1
type=tunnel
```

STEP 3 | Modify the strongSwan client's IPSec configuration file (**`ipsec.conf`**) and the IPSec password file (**`ipsec.secrets`**) to use recommended settings.

The `ipsec.secrets` file is usually found in the `/etc` folder.

© 2017 Palo Alto Networks, Inc.

Use the strongSwan client username as the certificate's common name.

Configure the following recommended settings in the `ipsec.conf` file:

```
conn <connection name>
keyexchange=ikev1
ikelifetime=1440m
keylife=60m
aggressive=yes
ike=aes-sha1-modp1024,aes256
esp=aes-sha1
xauth=client
left=<strongSwan/Linux-client-IP-address>
leftid=@#<hex of Group Name configured in the GlobalProtect gateway>
leftsourceip=%modeconfig
leftauth=psk
rightauth=psk
leftauth2=xauth
right=<gateway-IP-address>
rightsubnet=0.0.0.0/0
xauth_identity=<LDAP username>
auto=add
```

Configure the following recommended settings in the `ipsec.secrets` file:

```
: PSK <Group Password configured in the gateway>
 <username> : XAUTH "<user password>"
```

STEP 4 | Start strongSwan IPSec services and connect to the IPSec tunnel that you want the strongSwan client to use when authenticating to the GlobalProtect gateway.

- Ubuntu clients:

```
ipsec start
ipsec up <name>
```

- CentOS clients:

```
strongSwan start
strongswan up <name>
```

STEP 5 | Verify that the tunnel is set up correctly and the VPN connection is established to both the strongSwan client and the GlobalProtect gateway.

1. Verify the detailed status information on a specific connection (by naming the connection) or verify the status information for all connections from the strongSwan client:
   - Ubuntu clients:

```
ipsec statusall [<connection name>]
```

   - CentOS clients:

```
strongswan statusall [<connection name>]
```


© 2017 Palo Alto Networks, Inc.

2. Select **Network** > **GlobalProtect** > **Gateways.** Then, in the Info column, select **Remote Users** for the gateway configured for the connection to the strongSwan client. The strongSwan client should be listed under **Current Users**.

Enable Authentication Using Two-Factor Authentication

With two-factor authentication, the strongSwan client needs to successfully authenticate using both a certificate profile and an authentication profile to connect to the GlobalProtect gateway. The following workflow shows how to enable authentication for strongSwan clients using two-factor authentication.

STEP 1 | Set up the IPSec tunnel that the GlobalProtect gateway will use for communicating with a strongSwan client.

1. Select **Network** > **GlobalProtect** > **Gateways** and select the gateway name.
2. Select the **Certificate Profile** and **Authentication Profile** you want to use in the **Authentication** tab.
3. Select **Agent** > **Tunnel Settings** and specify the following settings to set up a tunnel:
   - Select the check box to **Enable X-Auth Support**.
   - If a **Group Name** and **Group Password** are already configured, remove them.
   - Click **OK** to save these tunnel settings.

STEP 2 | Verify that the default connection settings in the **`conn %default`** section of the IPSec tunnel configuration file (**`ipsec.conf`**) are correctly defined for the strongSwan client.

The `ipsec.conf` file usually resides in the `/etc` folder.



The configurations in this procedure are tested and verified for the following releases:

- *Ubuntu 14.0.4 with strongSwan 5.1.2 and CentOS 6.5 with strongSwan 5.1.3 for PAN-OS 6.1.*
- *Ubuntu 14.0.4 with strongSwan 5.2.1 for PAN-OS 7.0.*

Use the configurations in this procedure as a reference if you are using a different version of strongSwan. Refer to the strongSwan wiki for more information.

Configure the following recommended settings in the `ipsec.conf` file:

```
ikelifetime=20m
reauth=yes
rekey=yes
keylife=10m
rekeymargin=3m
rekeyfuzz=0%
keyingtries=1
type=tunnel
```

STEP 3 | Modify the strongSwan client's IPSec configuration file (**`ipsec.conf`**) and the IPSec password file (**`ipsec.secrets`**) to use recommended settings.

The `ipsec.secrets` file is usually found in the `/etc` folder.

Use the strongSwan client username as the certificate's common name.

Configure the following recommended settings in the `ipsec.conf` file:

```
conn <connection name>
keyexchange=ikev1
authby=xauthrsasig
```


© 2017 Palo Alto Networks, Inc.

```
ike=aes-sha1-modp1024
esp=aes-sha1
xauth=client
left=<strongSwan/Linux-client-IP-address>
leftcert=<client-certificate-without-password>
leftsourceip=%config
right=<GlobalProtect-gateway-IP-address>
rightid=%anyCN=<Subject-name-of-gateway-cert>"
rightsubnet=0.0.0.0/0
leftauth2=xauth
xauth_identity=<LDAP username>
auto=add
```

Configure the following recommended settings in the `ipsec.secrets` file:

```
<username> :XAUTH "<user password>"
::RSA <private key file> "<passphrase if used>"
```

STEP 4 | Start strongSwan IPSec services and connect to the IPSec tunnel that you want the strongSwan client to use when authenticating to the GlobalProtect gateway.

- Ubuntu clients:

```
ipsec start
ipsec up <name>
```

- CentOS clients:

```
strongSwan start
strongswan up <name>
```

STEP 5 | Verify that the tunnel is setup correctly and the VPN connection is established to both the strongSwan client and the GlobalProtect gateway.

1. Verify the detailed status information on a specific connection (by naming the connection) or verify the status information for all connections from the strongSwan client:
   - Ubuntu clients:

```
ipsec statusall [<connection name>]
```

   - CentOS clients:

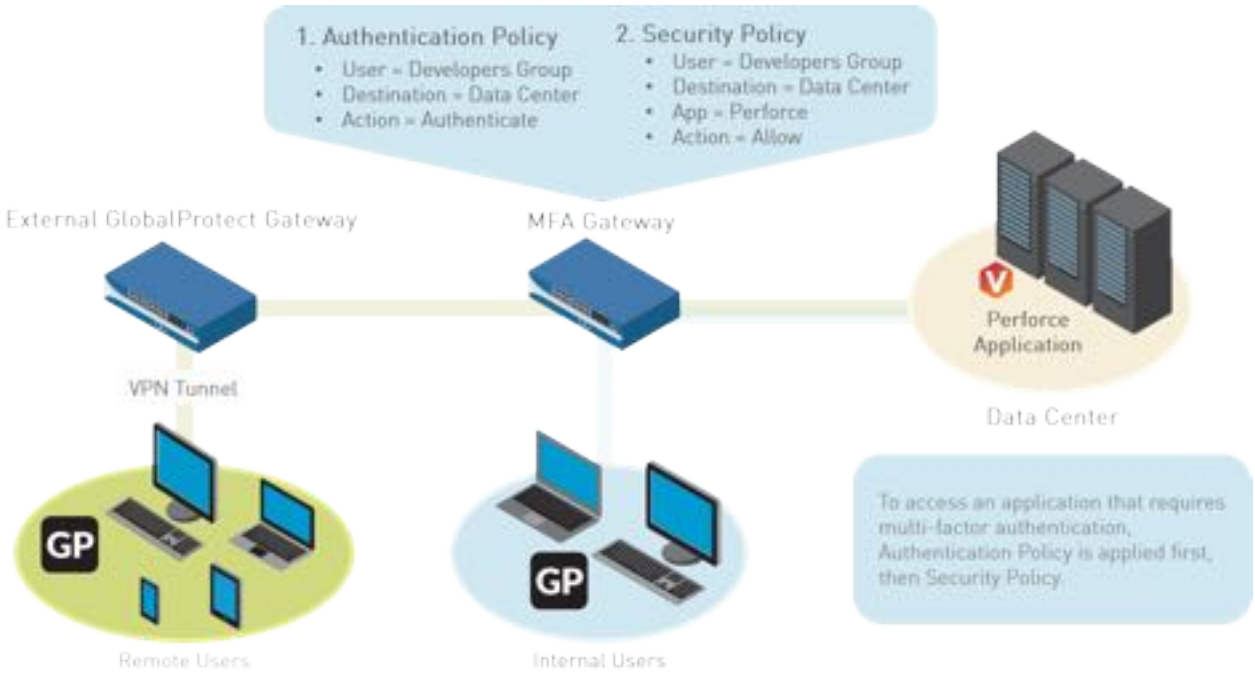
```
strongswan statusall [<connection name>]
```

2. Select **Network** > **GlobalProtect** > **Gateways**. Then, in the Info column, select **Remote Users** for the gateway configured for the connection to the strongSwan client. The strongSwan client should be listed under **Current Users**.

© 2017 Palo Alto Networks, Inc.

Configure GlobalProtect to Facilitate Multi-Factor Authentication Notifications

To protect critical applications and stop attackers from using stolen credentials to conduct lateral movement throughout your network, you can configure policy-based multi-factor authentication (MFA). This ensures that each user responds to multiple authentication challenges of different types (factors) before they can access highly sensitive services and applications.



If a user session matches the Authentication policy, the type of application or service determines the user experience for notifications about the authentication challenge:

- (Windows or Mac endpoints only) **Non-browser-based applications**—To facilitate MFA notifications for non-HTTP applications (such as Perforce) on Windows or Mac endpoints, a GlobalProtect client is required. When a session matches an Authentication policy rule, the firewall sends a UDP notification to the GlobalProtect client with an embedded URL link to the Authentication Portal page. The GlobalProtect client then displays this message as a pop up notification to the user.



- **Browser-based applications**—Browser-based applications do not require GlobalProtect to display notification messages to the user. When the firewall identifies a session as web-browsing traffic (based on App-ID), the firewall automatically presents the user with Authentication Portal page (previously called the Captive Portal page) specified in the Authentication policy rule. For more information, see Configure Multi-Factor Authentication.

© 2017 Palo Alto Networks, Inc.

To configure GlobalProtect to display MFA notifications for non-browser-based applications, use the following workflow:

STEP 1 | Before you configure GlobalProtect, configure multi-factor authentication on the firewall.



If you are using two-factor authentication with GlobalProtect to authenticate to the gateway or portal, a RADIUS server profile is required. If you are using GlobalProtect to notify the user about an authentication policy match (UDP message), a Multi Factor Authentication server profile is sufficient.

To use multi-factor authentication for protecting sensitive resources, the easiest solution is to integrate the firewall with an MFA vendor that is already established in your network. When your MFA structure is ready, you can start configuring the components of your authentication policy. For more information, refer to Configure Multi-Factor Authentication.

- Enable Captive Portal to record authentication timestamps and update user mappings.
- Create server profiles that define how the firewall will connect to the services that authenticate users.
- Assign the server profiles to an Authentication profile which specifies authentication parameters.
- Configure a Security policy rule that allows users to access the resources that require authentication.

STEP 2 | (External gateways only) For GlobalProtect to support multi-factor authentication on external gateways, you must Configure a response page for the ingress tunnel interface on the firewall:

1. Select **Device** > **Response Pages** > **MFA Login Page**.
2. Select and then **Export** the **Predefined** template to a location of your choice.
3. On your client system, use an HTML editor to customize the downloaded response page and save it with a unique filename.
4. Return to the **MFA Login Page** dialog on the firewall, **Import** your customized page, **Browse** to select the **Import File**, select the **Destination** (virtual system or shared location), click **OK**, and click **Close**.

STEP 3 | (External gateways only) Enable **Response Pages** as a permitted service on the **Interface Mgmt** profile:

1. Select **Network** > **Network Profiles** > **Interface Mgmt** and then select the profile.
2. In the **Permitted Services** area, select **Response Pages** and click **OK**.

STEP 4 | (External gateways only) Attach the **Interface Mgmt** profile to a tunnel interface:

1. Select **Network** > **Interfaces** > **Tunnel**, and the tunnel interface on which you want to use the response page.
2. Select **Advanced**, and then select the **Interface Mgmt** profile you configured in the previous step as the **Management Profile**.

STEP 5 | (External gateways only) **Enable User Identification** on the Zone associated with the tunnel interface (**Network** > **Zones** > **<tunnel-zone**).

STEP 6 | Configure GlobalProtect clients to support multi-factor authentication notifications for non-browser-based applications.

1. Select **Network** > **GlobalProtect** > **Portals** and select a portal configuration (or **Add** one).
2. Select **Agent** and then select an existing agent configuration or **Add** one.
3. In the **App** tab, specify the following:
   - Set **Enable Inbound Authentication Prompts from MFA Gateways** to **Yes**. To support multi-factor authentication (MFA), a GlobalProtect client must receive and acknowledge UDP prompts that are inbound from the gateway. Select **Yes** to enable a GlobalProtect client to receive and

© 2017 Palo Alto Networks, Inc.

acknowledge the prompt. By default, the value is set to **No** meaning GlobalProtect will block UDP prompts from the gateway.

- In **Network Port for Inbound Authentication Prompts (UDP)**, specify the port number a GlobalProtect client uses to receive inbound authentication prompts from MFA gateways. The default port is 4501. To change the port, specify a number from 1 to 65535.
- In **Trusted MFA Gateways**, specify the list of authentication gateways a GlobalProtect client will trust for multi-factor authentication. When a GlobalProtect client receives a UDP message on the specified network port, GlobalProtect displays an authentication message only if the UDP prompt comes from a trusted gateway.
- Configure the **Default Message for Inbound Authentication Prompts**. When users try to access a resource that requires additional authentication, GlobalProtect receives a UDP packet containing the inbound authentication prompt and displays this message. The UDP packet also contains the URL for the Authentication Portal page you specified in Configure Multi-Factor Authentication. GlobalProtect automatically appends the URL to the message. For example, to display the notification shown in the beginning of this topic enter the following message:

  ```
  You have attempted to access a protected resource that requires
  additional authentication. Proceed to authenticate at:
  ```

4. Save the agent configuration (click **OK** twice), and then **Commit** your changes.

© 2017 Palo Alto Networks, Inc.

Enable Delivery of GlobalProtect Client VSAs to a RADIUS Server

When communicating with GlobalProtect portals or gateways, GlobalProtect clients send information that includes the client IP address, operating system (OS), hostname, user domain, and GlobalProtect agent/app version. You can enable the firewall to send this information as Vendor-Specific Attributes (VSAs) to a RADIUS server during authentication (by default, the firewall does not send the VSAs). RADIUS administrators can then perform administrative tasks based on those VSAs. For example, RADIUS administrators might use the client OS attribute to define a policy that mandates regular password authentication for Microsoft Windows users and one-time password (OTP) authentication for Google Android users.

The following are prerequisites for this procedure:

- Import the Palo Alto Networks RADIUSdictionary into your RADIUS server.
- Configure a RADIUS server profile and assign it to an authentication profile: see Set Up External Authentication.
- Assign the authentication profile to a GlobalProtect portal or gateway: see Set Up Access to the GlobalProtect Portal or Configure a GlobalProtect Gateway.

STEP 1 | Log in to the firewall CLI.

STEP 2 | Enter the command for each VSA you want to send.

```
username@hostname> set authentication radius-vsa-on client-source-ip
username@hostname> set authentication radius-vsa-on client-os
username@hostname> set authentication radius-vsa-on client-hostname
username@hostname> set authentication radius-vsa-on user-domain
username@hostname> set authentication radius-vsa-on client-gp-version
```



If you later want to stop the firewall from sending particular VSAs, run the same commands but use the ***radius-vsa-off*** *option instead of* ***radius-vsa-on****.*

© 2017 Palo Alto Networks, Inc.

Enable Group Mapping

Because the agent or app running on your end-user systems requires the user to successfully authenticate before being granted access to GlobalProtect, the identity of each GlobalProtect user is known. However, if you want to be able to define GlobalProtect configurations and/or security policies based ongroup membership, the firewall must retrieve the list of groups and the corresponding list of members from your directory server. This is known as *group mapping*.

To enable this functionality, you must create an LDAP server profile that instructs the firewall how to connect and authenticate to the directory server and how to search the directory for the user and group information. After the firewall connects to the LDAP server and retrieves the group mappings, you can select groups when you define the agent configurations and security policies. The firewall supports a variety of LDAP directory servers, including Microsoft Active Directory (AD), Novell eDirectory, and Sun ONE Directory Server.

Use the following procedure to connect to your LDAP directory to enable the firewall to retrieve user-to-group mapping information:

STEP 1 | Create an LDAP Server Profile that specifies how to connect to the directory servers to which the firewall should connect to obtain group mapping information.

1. Select **Device** > **Server Profiles** > **LDAP** and click **Add**.
2. Enter a **Profile Name** to identify the server profile.
3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the **Location** where the profile is available.
4. For each LDAP server (up to four), **Add** and enter a **Name** (to identify the server), server IP address (**LDAP Server** field), and server **Port** (default 389).
5. Select the server **Type** from the drop-down: **active-directory**, **e-directory**, **sun**, or **other**.
6. If you want the device to use SSL or TLS for a more secure connection with the directory server, select the **Require SSL/TLS secured connection** check box (it is selected by default). The protocol that the device uses depends on the server **Port**:
   - 389 (default)—TLS (Specifically, the device uses the StartTLS operation, which upgrades the initial plaintext connection to TLS.)
   - 636—SSL
   - Any other port—The device first attempts to use TLS. If the directory server doesn't support TLS, the device falls back to SSL.
7. For additional security, you can select the **Verify Server Certificate for SSL sessions** check box (it is cleared by default) so that the device verifies the certificate that the directory server presents for SSL/TLS connections. To enable verification, you also have to select the **Require SSL/TLS secured connection** check box. For verification to succeed, the certificate must meet one of the following conditions:
   - It is in the list of device certificates: **Device** > **Certificate Management** > **Certificates** > **Device Certificates.** Import the certificate into the device, if necessary.
   - The certificate signer is in the list of trusted certificate authorities: **Device** > **Certificate Management** > **Certificates** > **Default Trusted Certificate Authorities**.
8. Click **OK**.

STEP 2 | Add the LDAP server profile to the User-ID Group Mapping configuration.

1. Select **Device** > **User Identification** > **Group Mapping Settings** and click **Add**.
2. Enter a **Name** for the configuration.
3. Select the **Server Profile** you just created.

© 2017 Palo Alto Networks, Inc.

4. Make sure the **Enabled** check box is selected.

STEP 3 | (**Optional**) Limit which groups can be selected in policy rules.

By default, if you don't specify groups, all groups are available in policy rules.

1. Add existing groups from the directory service:
    1. Select the **Group Include List** tab.
    2. In the Available Groups list, select the groups you want to appear in policy rules and click the Add icon .
2. If you want to base policy rules on user attributes that don't match existing user groups, create custom groups based on LDAP filters:
    1. Select the **Custom Group** tab and click **Add**.
    2. Enter a group **Name** that is unique in the group mapping configuration for the current firewall or virtual system. If the **Name** has the same value as the Distinguished Name (DN) of an existing AD group domain, the firewall uses the custom group in all references to that name (for example, in policies and logs).
    3. Specify an **LDAP Filter** of up to 2,048 UTF-8 characters, then click **OK**. The firewall doesn't validate LDAP filters.

    To optimize LDAP searches and minimize the performance impact on the LDAP directory server, use indexed attributes and reduce the search scope to include the user and group objects that you require for policy or visibility. Alternatively, you can create custom groups based on LDAP filters.

STEP 4 | Commit your changes.

Click **OK** and **Commit**.

GlobalProtect Gateways

- GlobalProtect Gateway Concepts
- Prerequisite Tasks for Configuring the GlobalProtect Gateway
- Configure a GlobalProtect Gateway

© 2017 Palo Alto Networks, Inc.

GlobalProtect Gateways Overview

Because the GlobalProtect™ configuration that the portal delivers to the agents includes the list of gateways the client can connect to, it is a good idea to configure the gateways before configuring the portal.

The GlobalProtect Gateways can be configured to provide two main functions:

- Enforce security policy for the GlobalProtect agents and apps that connect to it. You can also enable HIP collection on the gateway for enhanced security policy granularity. For more information on enabling HIP checks, see Host Information.
- Provide virtual private network (VPN) access to your internal network. VPN access is provided through an IPSec or SSL tunnel between the client and a tunnel interface on the gateway firewall.



You can also configure GlobalProtect gateways on VM-Series firewalls deployed in the AWS cloud. By deploying the VM-Series firewall in the AWS cloud you can quickly and easily deploy GlobalProtect gateways in any region without the expense or IT logistics that are typically required to set up this infrastructure using your own resources. For details, see Use Case: VM-Series Firewalls as GlobalProtect Gateways in AWS.

© 2017 Palo Alto Networks, Inc.

GlobalProtect Gateway Concepts

These sections provide information about gateway connection priority in a multiple gateway configuration and MIB support for GlobalProtect gateways.

- Types of Gateways on page 68
- Gateway Priority in a Multiple Gateway Configuration on page 68
- GlobalProtect MIB Support on page 69

Types of Gateways

GlobalProtect gateways provide security enforcement for traffic from GlobalProtect agents/apps. Additionally, if the HIP feature is enabled, the gateway generates a HIP report from the raw host data the clients submit and can use this information in policy enforcement.

You Configure a GlobalProtect Gateway on an interface on any Palo Alto Networks next-generation firewall. You can run both a gateway and a portal on the same firewall, or you can have multiple, distributed gateways throughout your enterprise.

You can configure any of the following types of gateways:

- **Internal**—An internal gateway is an interface on the internal network configured as a GlobalProtect gateway for applying security policy for access to internal resources. When used in conjunction with User-ID and/or HIP checks, an internal gateway can be used to provide a secure, accurate method of identifying and controlling traffic by user and/or device state. Internal gateways are useful in sensitive environments where authenticated access to critical resources is required. You can configure an internal gateway in either tunnel mode or non-tunnel mode. An agent connects to the internal gateway after performing internal host detection to determine the location of the endpoint.
- **External gateway (auto discovery)**—An external gateway resides outside of the corporate network and provides security enforcement and/or virtual private network (VPN) access for your remote users. The agent automatically connects to the external gateway depending on the priority you assign to the gateway, source region, and the response time (see Gateway Priority in a Multiple Gateway Configuration). When you configure an external gateway in the GlobalProtect portal agent configuration, auto discovery is the default. See Define the GlobalProtect Agent Configurations.
- **External gateway (manual)**—A manual external gateway also resides outside of the corporate network and provides security enforcement and/or VPN access for your remote users. The difference between the auto-discovery external gateway and the manual external gateway is that the GlobalProtect agent only connects to a manual external gateway when the user initiates a connection. You can also configure different authentication requirements for manual external gateways. To configure a manual gateway, you must identify the gateway as **Manual** in the GlobalProtect portal agent configuration. See Define the GlobalProtect Agent Configurations.

Gateway Priority in a Multiple Gateway Configuration

To enable secure access for your mobile workforce no matter where they are located, you can strategically deploy additional Palo Alto Networks next-generation firewalls and configure them as GlobalProtect gateways. To determine the preferred gateway to which your agents connect, add the gateways to a portal agent configuration and assign each gateway a connection priority. See Define the GlobalProtect Agent Configurations on page 86.

If a GlobalProtect portal agent configuration contains more than one gateway, the agent will attempt to communicate with all gateways listed in its agent configuration. The agent will then use priority and response time as to determine the gateway to which to connect. With GlobalProtect agent 4.0.2 and earlier releases, the agent connects to a lower priority gateway only if the response time for the higher priority gateway is greater than the average response time across all gateways.

© 2017 Palo Alto Networks, Inc.

For example, consider the following response times for gw1 and gw2:

Name	Priority	Response Time
gw1	Highest	80 ms
gw2	High	25 ms

The agent determines that the response time for the gateway with the highest priority (higher number) is greater than the average response time for both gateways (52.5 ms) and, as a result, connects to gw2. In this example, the agent did not connect to gw1 even though it had a higher priority because a response time of 80 ms was higher than the average for both.

Now consider the following response times for gw1, gw2, and a third gateway, gw3:

Name	Priority	Response Time
gw1	Highest	30 ms
gw2	High	25 ms
gw3	Medium	50 ms

In this example, the average response time for all gateways is 35 ms. The agent would then evaluate which gateways responded faster than the average response time and see that gw1 and gw2 both had faster response times. The agent would then connect to whichever gateway had the highest priority. In this example, the agent connects to gw1 because gw1 has the highest priority of all the gateways with response times below the average.

In addition to gateway priority, you add one or more source regions to an external gateway configuration, GlobalProtect recognizes the device region and only allows uses to connect to gateways that are configured for that region. For gateway choices, source region is considered first, then gateway priority.

In GlobalProtect agent 4.0.3 and later releases, the GlobalProtect agent prioritizes the gateways assigned highest, high, and medium priority ahead of gateways assigned a low or lowest priority regardless of response time. The GlobalProtect agent then appends any gateways assigned a low or lowest priority to the list of gateways. This ensures that the agent first attempts to connect to the gateways that you configure with a higher priority.

GlobalProtect MIB Support

Palo Alto Networks devices support standard and enterprise management information bases (MIBs) that enable you to monitor the device's physical state, utilization statistics, traps, and other useful information. Most MIBs use object groups to describe characteristics of the device using the Simple Network Management Protocol (SNMP) Framework. You must load these MIBs into your SNMP manager to monitor the objects (device statistics and traps) that are defined in the MIBs (for details, see Use an SNMP Manager to Explore MIBs and Objects in the PAN-OS 8.0 Administrator's Guide).

The PAN-COMMON-MIB—which is included with the enterprise MIBs—uses the panGlobalProtect object group. The following table describes the objects that make up the panGlobalProtect object group.

© 2017 Palo Alto Networks, Inc.

Object	Description
panGPGWUtilizationPct	Utilization (as a percentage) of the GlobalProtect gateway
panGPGWUtilizationMaxTunnels	Maximum number of tunnels allowed
panGPGWUtilizationActiveTunnels	Number of active tunnels

Use these SNMP objects to monitor utilization of GlobalProtect gateways and make changes as needed. For example, if the number of active tunnels reaches 80% or is higher than the maximum number of tunnels allowed, you should consider adding additional gateways.

© 2017 Palo Alto Networks, Inc.

Prerequisite Tasks for Configuring the GlobalProtect Gateway

Before you can configure the GlobalProtect gateway, you must have completed the following tasks:

- Created the interfaces (and zones) for the interface where you plan to configure each gateway. For gateways that require tunnel connections you must configure both the physical interface and the virtual tunnel interface. See Create Interfaces and Zones for GlobalProtect on page 15.
- Set up the gateway server certificates and SSL/TLS service profile required for the GlobalProtect agent to establish an SSL connection with the gateway. See Enable SSL Between GlobalProtect Components on page 17.
- Defined the authentication profiles and/or certificate profiles that will be used to authenticate GlobalProtect users. See Authentication on page 25.

© 2017 Palo Alto Networks, Inc.

Configure a GlobalProtect Gateway

After you have completed the prerequisite tasks, configure the GlobalProtect Gateways:

STEP 1 | Add a gateway.

1. Select **Network** > **GlobalProtect** > **Gateways** and click **Add**.
2. In the **General** screen, enter a **Name** for the gateway. The gateway name should have no spaces and, as a best practice, should include the location or other descriptive information to help users and administrators identify the gateway.
3. (**Optional**) Select the virtual system to which this gateway belongs from the **Location** field.

STEP 2 | Specify the network information that enables clients to connect to the gateway.

If you haven't created the network interface for the gateway, see Create Interfaces and Zones for GlobalProtect for instructions.



Do not attach an interface management profile that allows HTTP, HTTPS, Telnet, or SSH on the interface where you have configured a GlobalProtect portal or gateway because this enables access to your management interface from the Internet. Follow the Best Practices for Securing Administrative Access to ensure that you are securing administrative access to your firewalls in a way that will prevent successful attacks.

1. Select the **Interface** that clients will use for communication with the gateway.
2. Specify the **IP Address Type** and **IP address** for the gateway web service:
   - The IP address type can be **IPv4** (for IPv4 traffic only), **IPv6** (for IPv6 traffic only, or **IPv4 and IPv6.** Use **IPv4 and IPv6** if your network supports dual stack configurations, where IPv4 and IPv6 run at the same time.
   - The IP address must be compatible with the IP address type. For example, `172.16.1/0` for IPv4 addresses or `21DA:D3:0:2F3B` for IPv6 addresses. For dual stack configurations, enter both an IPv4 and IPv6 address.
3. Click **OK** to save changes.

STEP 3 | Specify how the gateway authenticates users.

If you haven't created an SSL/TLS service profile for the gateway, see Deploy Server Certificates to the GlobalProtect Components.

If you haven't set up the authentication profiles or certificate profiles, see Authentication for instructions.

Select **Authentication** and then configure any of the following:

- To secure communication between the gateway and the agents, select the **SSL/TLS Service Profile** for the gateway.

  To provide the strongest security, set the Min Version of the SSL/TLS service profile to TLSv1.2.
- To authenticate users with a local user database or an external authentication service, such as LDAP, Kerberos, TACACS+, SAML, or RADIUS (including OTP), **Add** a Client Authentication configuration with the following settings:
  - Enter a **Name** to identify the client authentication configuration.

© 2017 Palo Alto Networks, Inc.

- Identify the type of client to which this configuration applies. By default, the configuration applies to **Any** client, but you can customize the type of endpoint by **OS** (**Android**, **Chrome**, **iOS**, **Mac**, **Windows**, or **WindowsUWP**) or by third-party IPSec VPN clients (**X-Auth**).
- Select or add an **Authentication Profile** to authenticate an endpoint seeking access to the gateway.
- Enter an **Authentication Message** to help end users understand which credentials to use when logging in. The message can be up to 100 characters in length (default is `Enter login credentials`).
- To authenticate users based on a client certificate or a smart card/CAC, select the corresponding **Certificate Profile**.

- To use two-factor authentication, select both an authentication profile and a certificate profile. Keep in mind that the user must successfully authenticate using both methods to be granted access.

STEP 4 | Enable tunneling and configure the tunnel parameters.

The tunnel parameters are required if you are setting up an external gateway. If you are configuring an internal gateway, they are optional.

If you want to force use of SSL-VPN tunnel mode, clear the Enable IPSec check box. By default, SSL-VPN will only be used if the endpoint fails to establish an IPSec tunnel.

Extended authentication (X-Auth) is only supported on IPSec tunnels.

If you Enable X-Auth Support, GlobalProtect IPSec Crypto profiles are not applicable.

For more information on supported cryptographic algorithms, see Reference: GlobalProtect Agent Cryptographic Functions.

1. On the GlobalProtect Gateway Configuration dialog, select **Agent** > **Tunnel Settings.**
2. Select the **Tunnel Mode** check box to enable tunneling.
3. Select the **Tunnel Interface** you defined in 2 in Create Interfaces and Zones for GlobalProtect.
4. (**Optional**) Specify **Max User** for the maximum number of users that can access the gateway at the same time for authentication, HIP updates, and GlobalProtect agent updates (range varies based on the platform and is displayed when the field is empty).
5. Select a **GlobalProtect IPSec Crypto** profile to secure the VPN tunnels between GlobalProtect agents and gateways. The **default** profile uses AES-128-CBC encryption and sha1 authentication.

   You can also create a new IPSec crypto profile. To create a new profile, select **New GlobalProtect IPSec Crypto** in the same drop-down and configure the following:

   1. Enter a **Name** to identify the profile.
   2. **Add** the **Authentication** and **Encryption** algorithms that the VPN peers can use to negotiate the keys for securing the data in the tunnel:

      - **Encryption**—If you are not certain of what the VPN peers support, you can add multiple encryption algorithms in top-to-bottom order of most-to-least secure, as follows: **aes-256-gcm**, **aes-128-gcm**, **aes-128-cbc**. The peers negotiate the strongest algorithm to establish the tunnel.
      - **Authentication**—Select the authentication algorithm (**sha1**) to provide data integrity and authenticity protection. Although the authentication algorithm is required for the profile, this setting only applies to the AES-CBC cipher (**aes-128-cbc**). If you use an AES-GCM encryption algorithm (**aes-256-gcm** or **aes-128-gcm**), the setting is ignored because these ciphers natively provide ESP integrity protection.

© 2017 Palo Alto Networks, Inc.

3. Click **OK** to save the profile.

6. (**Optional**) Select **Enable X-Auth Support** if any endpoint needs to connect to the gateway by using a third-party VPN (for example, a VPNC client running on Linux). If you enable X-Auth, you must provide the **Group** name and **Group Password** if the endpoint requires it. By default, the user is not required to re-authenticate if the key used to establish the IPSec tunnel expires. To require users to re-authenticate, clear the option to **Skip Auth on IKE Rekey**.



Although X-Auth access is supported on iOS and Android endpoints, it provides limited GlobalProtect functionality on these endpoints. Instead, use the GlobalProtect app for simplified access to all the security features that GlobalProtect provides on iOS and Android endpoints. The GlobalProtect app for iOS is available at the Apple App Store. The GlobalProtect app for Android is available at Google Play.

STEP 5 | (**Optional**) Modify the default timeout settings for endpoints.

On the GlobalProtect Gateway Configuration dialog, select **Agent** > **Timeout Settings** and then configure the following settings:

- Modify the maximum **Login Lifetime** for a single gateway login session. The default login lifetime is 30 days—during the lifetime, the user stays logged in as long as the gateway receives a HIP check from the endpoint within the **Inactivity Logout** period. After this time, the login session automatically logs out.
- Modify the amount of time after which an inactive session is automatically logged out. The default **Inactivity Logout** period is 3 hours. A user is logged out of GlobalProtect if the gateway does not receive a HIP check from the endpoint during the configured amount of time.
- Modify the number of minutes after which idle users are logged out of GlobalProtect. The default period for **Disconnect on Idle** is 180 minutes. Users are logged out of GlobalProtect if the GlobalProtect agent has not routed traffic through the VPN tunnel in the configured amount of time. This setting applies to GlobalProtect agents that use the on-demand connect method only.

STEP 6 | (**Optional**) Configure authentication override settings to enable the gateway to generate and accept secure, encrypted cookies to authenticate the user. This capability allows the user to provide login credentials only once during a specified period of time (for example, every 24 hours).

By default, a gateway authenticates the user with an authentication profile and optional certificate profile. When authentication override is enabled, GlobalProtect caches the result of a successful login and uses the cookie to authenticate the user instead of prompting the user for credentials. For more information, see Cookie Authentication on the Portal or Gateway. If client certificates are required, the endpoint must also provide a valid certificate to be granted access.



In the event that you need to immediately block access to a device whose cookie has not yet expired (for example, if the device is lost or stolen), you can immediately Block Device Access *by adding the device to a block list.*

1. On the GlobalProtect Gateway Configuration dialog, select **Agent** > **Client Settings**.
2. **Add** a new agent configuration or select an existing configuration.
3. Enter a **Name** to identify the agent configuration.
4. Configure the following settings in the **Authentication Override** section:
   - **Generate cookie for authentication override**—Enable the gateway to generate encrypted, endpoint-specific cookies and issue the authentication cookies to the endpoint.
   - **Accept cookie for authentication override**—Enable the gateway to authenticate users with a valid, encrypted cookie. When the agent presents a valid cookie, the gateway verifies that the cookie was encrypted by the portal or gateway, decrypts the cookie, and then authenticates the user.

© 2017 Palo Alto Networks, Inc.

- **Cookie Lifetime**—Specify the hours, days, or weeks that the cookie is valid. Default is 24 hours. The range for hours is 1–72; for weeks, 1–52; and for days, 1–365. After the cookie expires, the user must enter login credentials, and the gateway subsequently encrypts a new cookie to send to the agent. This value can be the same as or different from the **Cookie Lifetime** you configure for the portal.
- **Certificate to Encrypt/Decrypt Cookie**—Select the RSA certificate to use to encrypt and decrypt the cookie. You must use the same certificate on the portal and gateways.

  As a best practice, configure the RSA certificate to use the strongest digest algorithm that your network supports.

  The portal and gateways use the RSA encrypt padding scheme PKCS#1 V1.5 to generate the cookie (using the public key of the certificate) and decrypt the cookie (using the private key of the certificate).

STEP 7 | Configure the user or user group and the endpoint OS to which the agent configuration applies.

The gateway uses the user/user group settings you specify to determine which configuration to deliver to the GlobalProtect agents that connect. Therefore, if you have multiple configurations, you must make sure to order them properly. As soon as the gateway finds a match, it will deliver the configuration. Therefore, more specific configurations must precede more general ones. See 10 for instructions on ordering the list of agent configurations.



Network settings are not required in internal gateway configurations in non-tunnel mode, because agents use the network settings assigned to the physical network adapter.

In a gateway agent configuration, select the **User/User Group** tab and configure the following settings:

- To deliver this configuration to agents or apps running on specific operating system, **Add** the OS (**Android**, **Chrome**, **iOS**, **Mac**, **Windows**, or **WindowsUWP**) to which this configuration applies. Or leave the value in this section set to **Any** to deploy the configuration based on user/group only.
- To restrict this configuration to a specific user and/or group, click **Add** in the User/User Group section of the window and then select the user or group you want to receive this configuration from the drop-down. Repeat this step for each user/group you want to add.

  Before you can restrict the configuration to specific groups, you must map users to groups as described in Enable Group Mapping.
- To restrict the configuration to users who have not yet logged in to their systems, select **pre-logon** from the User/User Group drop-down.
- To apply the configuration to any user regardless of login status (both pre-logon and logged in users), select **any** from the User/User Group drop-down.

STEP 8 | (**Tunnel Mode only**) Configure the ip pools available to assign to the virtual network adapter on the endpoint when an agent establishes a tunnel with the gateway.



IP pools and split tunnel settings are not required in internal gateway configurations in non-tunnel mode because agents use the network settings assigned to the physical network adapter.



You can optionally use address objects—which allow you to group specific source or destination addresses—when configuring gateway IP address pools or access routes.

© 2017 Palo Alto Networks, Inc.

In a gateway agent configuration, select **Agent** > **IP Pools** and configure any of the following settings and then click **OK**:

- To specify the authentication server IP address pool to assign addresses to endpoints that require static IP addresses, select the **Retrieve Framed-IP-Address attribute from authentication server** check box and then **Add** the subnet or IP address range to use to assign to remote users in the **Authentication Server IP Pool** area. When the tunnel is established, an interface is created on the remote user's computer with an address in the subnet or IP range that matches the Framed-IP attribute of the authentication server.

  The authentication server IP address pool must be large enough to support all concurrent connections. IP address assignment is static and is retained after the user disconnects.
- To specify the **IP Pool** to use to assign IP addresses, click **Add** and then specify the IP address range or address object to use. You can configure IPv6 or IPv4 addresses. As a best practice, use a different range of IP addresses from those assigned to endpoints that are physically connected to your LAN to ensure proper routing back to the gateway.

STEP 9 | (**Tunnel Mode only**) Configure the split tunnel settings to assign to the virtual network adapter on the endpoint when an agent establishes a tunnel with the gateway.

When configuring the access routes, keep in mind the following:

- More specific access routes take precedence over less specific routes.
- Avoid specifying the same access route as both an include and exclude access route as this leads to a misconfiguration.

To route only some traffic—likely traffic destined for your LAN—to GlobalProtect, specify the destination subnets or address object (of type **IP Netmask**) that must be included or excluded from the tunnel. In this case, traffic that is not destined for a specified access route will be routed through the endpoint's physical adapter rather than through the virtual adapter (the tunnel).

In a gateway agent configuration, select **Agent** > **Split Tunnel** and configure any of the following settings and then click **OK**:

- To disable split tunneling including direct access to local networks on Windows and Mac OS systems, enable **No direct access to local network**. In this case, users cannot send traffic to proxies or local resources while connected to GlobalProtect.
- To define what destination subnets to route through the tunnel click **Add** in the **Access Route** area and then enter the routes as follows:
  - (**Optional**) In the **Includes** area, **Add** the destination subnets or address object (of type IP Netmask) to route only some traffic—likely traffic destined for your LAN—to GlobalProtect. These are the routes the gateway pushes to the remote users' endpoint and thereby determines what traffic the users' endpoint can send through the VPN connection. You can include IPv6 or IPv4 subnets.

    The number of access routes the firewall supports varies by PAN-OS release version:

    - PAN-OS 8.0.0 and PAN-OS 8.0.1—Up to 100 include access routes, and, with GlobalProtect agent 4.0.2 or a later release, up to 200 include access routes
    - PAN-OS 8.0.2—Up to 100 include access routes and, with GlobalProtect agent 4.0.2 or a later release, up to 1000 include access routes
  - (**Optional**) In the **Excludes** area, **Add** the destination subnets or address object (of type IP Netmask) that you want the client to exclude. These routes will be sent through the endpoint's physical adapter rather than through the virtual adapter (the tunnel). Excluded routes should

© 2017 Palo Alto Networks, Inc.

be more specific than the included routes; otherwise, you may exclude more traffic than you intended. You can exclude IPv6 or IPv4 subnets. The firewall supports up to 100 exclude access routes, or with GlobalProtect agent 4.0.0 or a later release, up to 200 exclude access routes.

Excluding routes is not supported on Android. Only IPv4 routes are supported on Chrome.

STEP 10 | Arrange the gateway agent configurations so that the proper configuration is deployed to each agent.

When an agent connects, the gateway will compare the source information in the packet against the agent configurations you have defined. As with security rule evaluation, the gateway looks for a match starting from the top of the list. When it finds a match, it delivers the corresponding configuration to the agent or app.

- To move a gateway configuration up on the list of configurations, select the configuration and click **Move Up**.
- To move a gateway configuration down on the list of configurations, select the configuration and click **Move Down**.

STEP 11 | (**Tunnel Mode only**) Specify the network configuration settings for the endpoints.

Network settings are not required in internal gateway configurations in non-tunnel mode because in this case agents use the network settings assigned to the physical network adapter.

In a GlobalProtect Gateway Configuration, select the **Agent** > **Network Services** tab and configure the settings for endpoints in one of the following ways:

- If the firewall has an interface that is configured as a DHCP client, set the **Inheritance Source** to that interface and the GlobalProtect agent will be assigned the same settings received by the DHCP client. You can also **Inherit DNS Suffixes** from the inheritance source.
- Manually assign the DNS server(s) and suffix, and WINS servers by completing the corresponding fields.

STEP 12 | (**Optional**) Define the notification messages end users will see when a security rule with a host information profile (HIP) is enforced.

This step only applies if you have created host information profiles and added them to your security policies. For details on configuring the HIP feature and for more detailed information about creating HIP notification messages, see Host Information.

In a GlobalProtect Gateway Configuration, select the **Agent** > **HIP Notification** tab and **Add** a new HIP Notification configuration:

1. From the **Host Information** drop-down, select the HIP object or profile to which this message applies.
2. Select **Match Message** or **Not Match Message** and then **Enable** notifications, depending on whether you want to display the message when the corresponding HIP profile is matched in policy or when it is not matched. In some cases, you might want to create messages for both a match and a non-match, depending on the objects on which you are matching and what your objectives are for the policy. For the Match Message, you can also enable the option to **Include Mobile App List** to indicate what applications can trigger the HIP match.
3. Select whether you want to display the message as a **System Tray Balloon** or as a **Pop Up Message**.
4. Enter and format the text of your message in the Template text box and then click **OK**.
5. Repeat these steps for each message you want to define.

STEP 13 | Save the gateway configuration.

© 2017 Palo Alto Networks, Inc.

1. Click **OK** to save the settings and close the GlobalProtect Gateway Configuration dialog.
2. **Commit** the changes.

GlobalProtect Portals

- GlobalProtect Portal Overview on page 81
- Prerequisite Tasks for Configuring the GlobalProtect Portal on page 82
- Set Up Access to the GlobalProtect Portal on page 83
- Define the GlobalProtect Agent Configurations on page 86
- Customize the GlobalProtect Agent on page 91
- Customize the GlobalProtect Portal Login, Welcome, and Help Pages on page 99
- GlobalProtect Clientless VPN on page 135

© 2017 Palo Alto Networks, Inc.

GlobalProtect Portal Overview

The GlobalProtect Portal provides the management functions for your GlobalProtect infrastructure. Every endpoint that participates in the GlobalProtect network receives configuration information from the portal, including information about available gateways as well as any client certificates that may be required to connect to the gateways. In addition, the portal controls the behavior and distribution of the GlobalProtect agent software to both Mac and Windows laptops.



The portal does not distribute the GlobalProtect app for use on mobile devices. To get the GlobalProtect app for mobile devices, end users must download it from the store for their device: App Store for iOS, Google Play for Android, Chrome Web Store for Chromebooks, or Microsoft Store for Windows 10 UWP. However, the agent configurations that get deployed to mobile app users does control the gateway(s) to which the mobile devices have access. For more details on supported versions, see What Client OS Versions are Supported with GlobalProtect?

In addition to distributing GlobalProtect client software, you can configure the GlobalProtect portal to provide secure remote access to common enterprise web applications that use HTML, HTML5, and Javascript technologies. Users have the advantage of secure access from SSL-enabled web browsers without installing GlobalProtect client software. This is useful when you need to enable partner or contractor access to applications, and to safely enable unmanaged assets, including personal devices. Refer to GlobalProtect Clientless VPN.

© 2017 Palo Alto Networks, Inc.

Prerequisite Tasks for Configuring the GlobalProtect Portal

Before you can configure the GlobalProtect portal, you must complete the following tasks:

- Create the interfaces (and zones) for the firewall interface where you plan to configure the portal. See Create Interfaces and Zones for GlobalProtect.
- Set up the portal server certificate, gateway server certificate, SSL/TLS service profiles, and, optionally, any client certificates to deploy to end users to enable SSL/TLS connections for the GlobalProtect™ services. See Enable SSL Between GlobalProtect Components.
- Define the optional authentication profiles and certificate profiles that the portal can use to authenticate GlobalProtect users. See Authentication.
- Configure a GlobalProtect Gateway and understand Gateway Priority in a Multiple Gateway Configuration.

© 2017 Palo Alto Networks, Inc.

Set Up Access to the GlobalProtect Portal

After you have completed the Prerequisite Tasks for Configuring the GlobalProtect Portal, configure the GlobalProtect portal as follows:

STEP 1 | Add the portal.

1. Select **Network** > **GlobalProtect** > **Portals** and click **Add**.
2. On the **General** page, enter a **Name** for the portal. The name cannot contain spaces.
3. (**Optional**) Select the virtual system to which this portal belongs from the **Location** field.

STEP 2 | Specify network settings to enable agents to communicate with the portal.

If you have not yet created the network interface for the portal, see Create Interfaces and Zones for GlobalProtect for instructions. If you have not yet created an SSL/TLS service profile for the portal, see Deploy Server Certificates to the GlobalProtect Components.



Do not attach an interface management profile that allows HTTP, HTTPS, Telnet, or SSH on the interface where you have configured a GlobalProtect portal or gateway because this enables access to your management interface from the Internet. Follow the Best Practices for Securing Administrative Access to ensure that you are securing administrative access to your firewalls in a way that will prevent successful attacks.

1. Select the **Interface**.
2. Specify the **IP Address Type** and **IP address** for the portal web service:
   - The IP address type can be **IPv4** (for IPv4 traffic only), **IPv6** (for IPv6 traffic only, or **IPv4 and IPv6.** Use **IPv4 and IPv6** if your network supports dual stack configurations, where IPv4 and IPv6 run at the same time.
   - The IP address must be compatible with the IP address type. For example, `172.16.1/0` for IPv4 addresses or `21DA:D3:0:2F3B` for IPv6 addresses. For dual stack configurations, enter both an IPv4 and IPv6 address.
3. Select an **SSL/TLS Service Profile**.

STEP 3 | Disable the login page entirely or choose your own login page or help page. Although optional, a custom login or help page lets you decide on the look and content of the pages. See Customize the GlobalProtect Portal Login, Welcome, and Help Pages.

- Choose a **Portal Login Page** for user access to the portal or import a new one, or **Disable** access to the GlobalProtect portal login page from a web browser.
- Choose a **App Help Page** to assist the user with GlobalProtect or import a new one.

STEP 4 | Specify how the portal authenticates the users.

If you have not yet created a server certificate for the portal and issued gateway certificates, see Deploy Server Certificates to the GlobalProtect Components.

On the GlobalProtect Portal Configuration dialog, select **Authentication**, and then configure any of the following:

- To secure communication between the portal and the agents, select the **SSL/TLS Service Profile** you configured for the portal.
- To authenticate users using a local user database or an external authentication service, such as LDAP, Kerberos, TACACS+, SAML, or RADIUS (including OTP), Define the GlobalProtect Client Authentication Configurations.

© 2017 Palo Alto Networks, Inc.

STEP 5 | Save the portal configuration.

1. Click **OK** to save the settings and close the GlobalProtect Portal Configuration dialog.
2. **Commit** the changes.

© 2017 Palo Alto Networks, Inc.

Define the GlobalProtect Client Authentication Configurations

Each GlobalProtect client authentication configuration specifies the settings that enable the user to authenticate with the GlobalProtect portal. You can customize the settings for each OS or you can configure the settings to apply to all devices. For example, you can configure Android users to use RADIUS authentication and Windows users to use LDAP authentication. You can also customize the client authentication for users who access the portal from a web browser (to download the GlobalProtect agent) or for third-party IPSec VPN (X-Auth) access to GlobalProtect gateways.

STEP 1 | Set Up Access to the GlobalProtect Portal on page 83.

1. Select **Network** > **GlobalProtect** > **Portals**.
2. Select the portal configuration to which you are adding the client configuration and then select the **Authentication** tab.

STEP 2 | Specify how the portal authenticates the users.

You can configure the GlobalProtect portal to authenticate users using a local user database or an external authentication service, such as LDAP, Kerberos, TACACS+, SAML, or RADIUS (including OTP). If you have not yet set up the authentication profiles and/or certificate profiles, see Authentication on page 25 for instructions.

In the Client Authentication area, **Add** a new configuration with the following settings:

- Enter a **Name** to identify the client authentication configuration.
- Specify the endpoints to which to deploy this configuration. By default, the configuration applies to all endpoints. Otherwise, you can apply the configuration to endpoints running a specific **OS** (**Android**, **Chrome**, **iOS**, **Mac**, **Windows**, or **WindowsUWP**) or to endpoints that access the portal from a web **Browser** with the intent of downloading the GlobalProtect agent or to create a new client authentication specifically for GlobalProtect Clientless VPN on page 135.
- Select or add an **Authentication Profile** for authenticating an endpoint that tries to access the gateway.
- Enter an **Authentication Message** to help end users understand which credentials to use when logging in. The message can be up to 100 characters in length (default is `Enter login credentials`).

STEP 3 | Arrange the client authentication configurations with OS-specific configurations at the top of the list, and configurations that apply to **Any** OS at the bottom of the list. As with security rule evaluation, the portal looks for a match starting from the top of the list. When it finds a match, it delivers the corresponding configuration to the agent or app.

- To move a client authentication configuration up on the list of configurations, select the configuration and click **Move Up**.
- To move a client authentication configuration down on the list of configurations, select the configuration and click **Move Down**.

STEP 4 | (**Optional**) To enable two-factor authentication using an authentication profile and a certificate profile, configure both in this portal configuration.

Keep in mind the portal must authenticate the client by using both methods before the user can gain access.

© 2017 Palo Alto Networks, Inc.

Select the corresponding **Certificate Profile** to authenticate users based on a client certificate or smart card.

The Common Name (CN) and, if applicable, the Subject Alternative Name (SAN) fields of the certificate must exactly match the IP address or FQDN of the interface where you configure the portal or HTTPS connections to the portal will fail.

STEP 5 | Save the portal configuration.

1. Click **OK** to save the settings and close the GlobalProtect Portal Configuration dialog.
2. **Commit** the changes.

Define the GlobalProtect Agent Configurations

After a GlobalProtect user connects to the portal and is authenticated by the GlobalProtect portal, the portal sends the agent configuration to the agent or app, based on the settings you defined. If you have different roles for users or groups that need specific configurations, you can create a separate agent configuration for each user type or user group. The portal uses the OS of the endpoint and the username or group name to determine the agent configuration to deploy. As with other security rule evaluations, the portal starts to search for a match at the top of the list. When it finds a match, the portal sends the right configuration to the agent or app.

The configuration can include the following:

- A list of gateways to which the client can connect.
- Among the external gateways, any gateway that the user can manually select for the session.
- The root CA certificate required to enable the agent or app to establish an SSL connection with the GlobalProtect gateway(s).
- The root CA certificate for SSL forward proxy decryption.
- The client certificate that the endpoint should present to the gateway when it connects. This configuration is required only if mutual authentication between the client and the portal or gateway is required.
- A secure encrypted cookie that the endpoint should present to the portal or gateway when it connects. The cookie is included only if you enable the portal to generate one.
- The settings the endpoint uses to determine whether it is connected to the local network or to an external network.
- Settings for the behavior of the agent or app, such as what the end users can see in their display, whether they can save their GlobalProtect password, and whether they are prompted to upgrade their software.

If the portal is down or unreachable, the agent will use the cached version of its agent configuration from its last successful portal connection to obtain settings, including the gateway(s) to which the agent can connect, what root CA certificate(s) to use to establish secure communication with the gateway(s), and what connect method to use.

Use the following procedure to create an agent configuration.

STEP 1 | Add one or more trusted root CA certificates to the portal agent configuration to enable the GlobalProtect client to verify the identity of the portal and gateways.

The portal deploys the certificate in a certificate file which is read only by GlobalProtect.

1. Select **Network** > **GlobalProtect** > **Portals**.

© 2017 Palo Alto Networks, Inc.

2. Select the portal configuration to which you are adding the agent configuration and then select the **Agent** tab.
3. In the **Trusted Root CA** field, **Add** and then select the CA certificate that was used to issue the gateway and/or portal server certificates.

   The web interface presents a list of CA certificates that are imported on the firewall serving as the GlobalProtect portal. The web interface also excludes end-entity certificates, sometimes referred to as leaf certificates, from the list of certificates you can select. You can also **Import** a new CA certificate.

   Use the following best practices when creating and adding certificates:

   - *Use the same certificate issuer to issue certificates for all of your gateways.*
   - *Add the entire certificate chain (trusted root CA and intermediate CA certificates) to the portal agent configuration.*
4. (Optional) Deploy additional CA certificates for purposes other than GlobalProtect (for example, SSL forward proxy decryption).

   This option enables you to use the portal to deploy certificates to the endpoint and the agent to install them in the local root certificate store. This can be useful if you do not have another method for distributing these server certificates or prefer to use the portal for certificate distribution.

   For SSL forward proxy decryption, you specify the forward trust certificate the firewall uses (on Windows and Mac endpoints only) to terminate the HTTPS connection, inspect the traffic for policy compliance, and re-establish the HTTPS connection to forward the encrypted traffic.

   1. Add the certificate as described in the previous step.
   2. To the right of the certificate, select **Install in Local Root Certificate Store**.

      The portal automatically sends the certificate when the user logs in to the portal and installs it in the client's local store thus eliminating the need for you to install the certificate manually.

STEP 2 | Add an agent configuration.

The agent configuration specifies the GlobalProtect configuration settings to deploy to the connecting agents/apps. You must define at least one agent configuration.

1. In the Agent area, **Add** a new configuration.
2. Enter a **Name** to identify the configuration. If you plan to create multiple configurations, make sure the name you define for each is descriptive enough to allow you to distinguish them.

STEP 3 | (**Optional**) Configure settings to specify how users with this configuration will authenticate with the portal.

If the gateway is to authenticate the clients by using a client certificate, you must select the source that distributes the certificate.

On the **Authentication** tab, configure any of the following authentication settings:

- To enable users to authenticate with the portal using client certificates, select the **Client Certificate** source (**SCEP**, **Local**, or **None**) that distributes the certificate and its private key to an endpoint. If you use an internal CA to distribute certificates to clients, select **None** (**default**). To enable the portal to generate and send a machine certificate to the agent for storage in the local certificate store and use the certificate for portal and gateway authentication, select **SCEP** and the associated SCEP profile. These certificates are device-specific and can only be used on the endpoint to which it was issued. To use the same certificate for all endpoints, select a certificate that is **Local** to the portal. With **None**, the portal does not push a certificate to the client, but you can use can other ways to get a certificate to the client's endpoint.
- Specify whether to **Save User Credentials**. Select **Yes** to save the username and password (default), **Save Username Only** to save only the username, or **No** to never save credentials.

© 2017 Palo Alto Networks, Inc.

If you configure the portal or gateways to prompt for a dynamic password such as a one-time password (OTP), the user must enter a new password at each login. In this case, the GlobalProtect agent/app ignores the selection to save both the username and password, if specified, and saves only the username. For more information, see Enable Two-Factor Authentication Using One-Time Passwords (OTPs).

STEP 4 | If the GlobalProtect endpoint does not require tunnel connections when it is on the internal network, configure internal host detection.

1. On the **Internal** tab, select the **Internal Host Detection** check box.
2. Enter the **IP Address** of a host that can be reached from the internal network only. You can configure **IPv4** or **IPv6** addressing for **Internal Host Detection**. The IP address you specify must be compatible with the IP address type. For example, 172.16.1.0 for IPv4 or 21DA:D3:0:2F3b for IPv6.
3. Enter the DNS **Hostname** for the IP address you entered. Clients that try to connect to GlobalProtect attempt to do a reverse DNS lookup on the specified address. If the lookup fails, the client determines that it is on the external network and then initiates a tunnel connection to a gateway on its list of external gateways.

STEP 5 | Set up access to a third-party mobile endpoint management system.

This step is required if the mobile devices using this configuration will be managed by a third-party mobile endpoint management system. All devices will initially connect to the portal and, if a third-party mobile endpoint management system is configured on the corresponding portal agent configuration, the device will be redirected to it for enrollment.

1. Enter the IP address or FQDN of the device check-in interface associated with your mobile endpoint management system. The value you enter here must exactly match the value of the server certificate associated with the device check-in interface. You can specify an IPv6 or IPv4 address.
2. Specify the **Enrollment Port** on which the mobile endpoint management system will be listening for enrollment requests. This value must match the value set on the mobile endpoint management system (default=443).

STEP 6 | Configure the user or user group and the endpoint OS to which the agent configuration applies.

The portal uses the user/user group settings you specify to determine which configuration to deliver to the GlobalProtect agents that connect. Therefore, if you have multiple configurations, you must make sure to order them properly. As soon as the portal finds a match, it will deliver the configuration. Therefore, more specific configurations must precede more general ones. See 12 for instructions on ordering the list of agent configurations.

Select the **User/User Group** tab and then specify any users, user groups, and/or operating systems to which this configuration should apply:

- To deliver this configuration to agents or apps running on specific operating system, **Add** the OS (**Android**, **Chrome**, **iOS**, **Mac**, **Windows**, or **WindowsUWP**) to which this configuration applies. Or leave the value in this section set to **Any** to deploy the configuration based on user/group only.
- To restrict this configuration to a specific user and/or group, click **Add** in the User/User Group section of the window and then select the user or group you want to receive this configuration from the drop-down. Repeat this step for each user/group you want to add.



Before you can restrict the configuration to specific groups, you must map users to groups as described in Enable Group Mapping*.*

- To restrict the configuration to users who have not yet logged in to their systems, select **pre-logon** from the User/User Group drop-down.

© 2017 Palo Alto Networks, Inc.

- To apply the configuration to any user regardless of login status (both pre-logon and logged in users), select **any** from the User/User Group drop-down.

STEP 7 | Specify the external gateways to which users with this configuration can connect.



Consider the following best practices when you configure the gateways:

- *If you are adding both internal and external gateways to the same configuration, make sure to enable Internal Host Detection. See 4 in* Define the GlobalProtect Agent Configurations *for instructions.*
- *To learn more about how a GlobalProtect client determines the gateway to which it should connect, see* Gateway Priority in a Multiple Gateway Configuration.

1. Click **Add** on the **External** tab.
2. Enter a descriptive **Name** for the gateway. The name you enter here should match the name you defined when you configured the gateway and should be descriptive enough for users to know the location of the gateway they are connected to.
3. Enter the FQDN or IP address of the interface where the gateway is configured in the **Address** field. You can configure an IPv4 or IPv6 address. The address you specify must exactly match the Common Name (CN) in the gateway server certificate.
4. **Add** one or more **Source Regions** for the gateway, or select **Any** to make the gateway available to all regions. When users connect, GlobalProtect recognizes the device region and only allows uses to connect to gateways that are configured for that region. For gateway choices, source region is considered first, then gateway priority.
5. Set the **Priority** of the gateway by clicking in the field and selecting a value:
   - If you have only one external gateway, you can leave the value set to **Highest** (the default).
   - If you have multiple external gateways, you can modify the priority values (ranging from **Highest** to **Lowest**) to indicate a preference for the specific user group to which this configuration applies. For example, if you prefer that the user group connects to a local gateway you would set the priority higher than that of more geographically distant gateways. The priority value is then used to weight the agent's gateway selection algorithm.
   - If you do not want agents to automatically establish tunnel connections with the gateway, select **Manual only**. This setting is useful in testing environments.
6. Select the **Manual** check box if you want to allow users to be able to manually switch to the gateway.

STEP 8 | Specify the internal gateways to which users with this configuration can connect.



Make sure you do not use on-demand as the connect method if your configuration includes internal gateways.

1. On the **Internal** tab, click **Add** in the **Internal Gateways** section.
2. Enter a descriptive **Name** for the gateway. The name you enter here should match the name you defined when you configured the gateway and should be descriptive enough for users to know the location of the gateway they are connected to.
3. Enter the FQDN or IP address of the interface where the gateway is configured in the **Address** field. You can configure an IPv4 or IPv6 address. The address you specify must exactly match the Common Name (CN) in the gateway server certificate.
4. (**Optional**) **Add** one or more **Source Addresses** to the gateway configuration. The source address can be an IP subnet or range. It can also be a predefined address. GlobalProtect supports both IPv6 and IPv4 addresses. When users connect, GlobalProtect recognizes the source address of the device and only allows users to connect to gateways that are configured for that address.
5. Click **OK** to save your changes.

© 2017 Palo Alto Networks, Inc.

6. (**Optional**) **Add** a **DHCP Option 43 Code** to the gateway configuration. You can include one or more sub-option codes associated with the vendor-specific information (Option 43) that the DHCP server has been configured to offer the client. For example, you might have a sub-option code 100 that is associated with an IP address of 192.168.3.1.

   When a user connects, the GlobalProtect portal sends the list of option codes in the portal configuration to the GlobalProtect agent and the agent selects gateways indicated by the options.

   When both the source address and DHCP options are configured, the list of available gateways presented to the client is based on the combination (union) of the two configurations.

   DHCP options are supported on Windows and Mac endpoints only. DHCP options cannot be used to select gateways that use IPv6 addressing.

7. (**Optional**) Select **Internal Host Detection** to allow the GlobalProtect agent to determine if it is inside the enterprise network. When the user attempts to log in, the agent does a reverse DNS lookup of the internal **Hostname** to the specified **IP Address**.

   The host serves as a reference point that is reachable if the endpoint is inside the enterprise network. If the agent finds the host, the endpoint is inside the network and the agent connects to an internal gateway; if the agent fails to find the internal host, the endpoint is outside the network and the agent establishes a tunnel to one of the external gateways.

   You can configure **IPv4** or **IPv6** addressing for **Internal Host Detection**. The IP address you specify must be compatible with the IP address type. For example, 172.16.1.0 for IPv4 or 21DA:D3:0:2F3b for IPv6.

STEP 9 | Customize the behavior of the GlobalProtect agent for users with this configuration.

Select the **App** tab and then modify the agent settings as desired. For more details about each option, see Customize the GlobalProtect Agent.

STEP 10 | (**Optional**) Define any custom host information profile (HIP) data that you want the agent to collect and/or exclude HIP categories from collection.

This step only applies if you plan to use the HIP feature and there is information you want to collect that cannot be collected using the standard HIP objects or if there is HIP information that you are not interested in collecting. See Host Information for details on setting up and using the HIP feature.

1. Select **Data Collection** and enable the GlobalProtect agent to **Collect HIP Data**.
2. Select **Exclude Categories** to exclude specific categories and/or vendors, applications, or versions within a category. For more details, see 3 in Configure HIP-Based Policy Enforcement.
3. Select **Custom Checks** to define any custom data you want to collect from hosts running this agent configuration, and add the category and vendor. For more details, see 2 in Host Information.

STEP 11 | Save the agent configuration.

1. Click **OK** to save the settings and close the Configs dialog.
2. If you want to add another agent configuration, repeat 2 through 11.

STEP 12 | Arrange the agent configurations so that the proper configuration is deployed to each agent.

When an agent connects, the portal will compare the source information in the packet against the agent configurations you have defined. As with security rule evaluation, the portal looks for a match starting from the top of the list. When it finds a match, it delivers the corresponding configuration to the agent or app.

- To move an agent configuration up on the list of configurations, select the configuration and click **Move Up**.

© 2017 Palo Alto Networks, Inc.

- To move an agent configuration down on the list of configurations, select the configuration and click **Move Down**.

STEP 13 | Save the portal configuration.

1. Click **OK** to save the settings and close the GlobalProtect Portal Configuration dialog.
2. **Commit** the changes.

Customize the GlobalProtect Agent

The portal agent configuration allows you to customize how your end users interact with the GlobalProtect agents installed on their systems or the GlobalProtect app installed on their mobile devices. You can define different agent settings for the different GlobalProtect agent configurations you create. For more information on GlobalProtect client requirements, see What Client OS Versions are Supported with GlobalProtect?

You can customize the display and behavior of the agent. For example, you can specify the following:

- What menus and views users can access.
- Whether users can disable the agent (applies to the user-logon connect method only).
- Whether to display a welcome page upon successful login. You can also configure whether or not the user can dismiss the welcome page and you can create custom welcome and help pages that explain how to use GlobalProtect within your environment. See Customize the GlobalProtect Portal Login, Welcome, and Help Pages.
- Whether agent upgrades occur automatically or whether users are prompted to upgrade.
- Prompt users if multi-factor authentication is needed to access sensitive network resources.



You can also define agent settings directly from the Windows registry or the global Mac plist. For Windows clients you can also define agent settings directly from the Windows installer (Msiexec). Settings defined in the portal agent configurations in the web interface take precedence over settings defined in the Windows registry/Msiexec or the Mac plist. For more details, see Deploy Agent Settings Transparently.

Additional options that are available through the Windows command line (Msiexec) or Windows registry only, enable you to (for more information, see Customizable Agent Settings):

- Specify whether the agent should prompt the end user for credentials if Windows SSO fails.
- Specify the default portal IP address (or hostname).
- Enable GlobalProtect to initiate a VPN connection before the user logs into the endpoint.
- Deploy scripts that run before or after GlobalProtect establishes a VPN connection or after GlobalProtect disconnects the VPN connection.
- Enable the GlobalProtect agent to wrap third-party credentials on the Windows client, allowing for SSO when using a third-party credential provider.

Use the following procedure to customize the GlobalProtect agent.

STEP 1 | Select the **Agent** tab in the agent configuration you want to customize.



You can also configure most settings that are on the App tab from a group policy by adding settings to the Windows registry/Mac plist. On Windows systems, you can also set them using the Msiexec utility on the command line during the agent installation. However, settings defined in the web interface or the CLI take precedence over registry/plist settings. See Deploy Agent Settings Transparently *for details.*

© 2017 Palo Alto Networks, Inc.

1. Select **Network** > **GlobalProtect** > **Portals** and select the portal configuration for which you want to add an agent configuration (or **Add** a new configuration).
2. Select the **Agent** tab and select the configuration you want to modify (or **Add** a new configuration).
3. Select the **App** tab.

   The App Configurations area displays the options with default values that you can customize for each agent configuration. When you change the default behavior, the web interface changes the color from gray to the default text color.

STEP 2 | Specify the **Connect Method** that an agent or app uses for its GlobalProtect connection.



*Use the **Pre-logon (Always On)**, **Pre-logon then On-demand**, or **User-log on (Always On)** connect method to access the network using an internal gateway.*

In the App Configurations area, configure any of the following options:

- Select a **Connect Method**:
- **User-logon (Always On)**—The GlobalProtect agent automatically connects to the portal as soon as the user logs in to the endpoint (or domain). When used in conjunction with SSO (Windows users only), GlobalProtect login is transparent to the end user.

  On iOS endpoints, this setting prevents one-time password (OTP) applications from working because GlobalProtect forces all traffic to go through the tunnel.
- **Pre-logon (Always On)**—Authenticates the user and establishes a VPN tunnel to the GlobalProtect gateway before the user logs in to the client. This option requires that you use an external PKI solution to pre-deploy a machine certificate to each endpoint that receives this configuration. See Remote Access VPN with Pre-Logon for details about pre-logon.
- **On-demand (Manual user initiated connection)**—Users will have to manually launch the agent to connect to GlobalProtect. Use this connect method for external gateways only.
- **Pre-logon then On-demand**—Similar to the **Pre-logon (Always On)** connect method, this connect method (which requires Content Release version 590-3397 or later) enables the GlobalProtect agent to authenticate the user and establish a VPN tunnel to the GlobalProtect gateway before the user logs in to the client. Unlike the pre-logon connection method, after the user logs in to the client, users must manually launch the agent to connect to GlobalProtect if the connection is terminated for any reason. The benefit of this option is that you can allow a user to specify a new password after password expiration or a user forgets their password but still require the user to manually initiate the connection after the user logs in.

STEP 3 | Specify whether to enforce GlobalProtect connections for network access.



To enforce GlobalProtect for network access, we recommend that you enable this feature only for users that connect in User-logon or Pre-logon modes. Users that connect in On-demand mode may not be able to establish a connection within the permitted grace periods.

In the App Configurations area, configure any of the following options:

- To force all network traffic to traverse a GlobalProtect tunnel, set **Enforce GlobalProtect Connection for Network Access** to **Yes**. By default, GlobalProtect is not required for network access meaning users can still access the internet if GlobalProtect is disabled or disconnected. To provide instructions to users before traffic is blocked, configure a **Traffic Blocking Notification Message** and optionally specify when to display the message (**Traffic Blocking Notification Delay**).

© 2017 Palo Alto Networks, Inc.



*When **Enforce GlobalProtect Connection for Network Access** is enabled, you may want to consider allowing users to disable the GlobalProtect agent with a passcode. The **Enforce GlobalProtect Connection for Network Access** feature enhances the network security by requiring a GlobalProtect VPN connection for network access. On rare occasions, devices may fail to connect to the VPN and require remote administrative login for troubleshooting. By disabling the GlobalProtect agent (for Windows or Mac) using the passcode provided by the administrator during the troubleshooting session, you can allow administrators to connect to your device remotely.*

- To permit traffic required to establish a connection with a captive portal, specify a **Captive Portal Exception Timeout**. The user must authenticate with the portal before the timeout expires. To provide additional instructions, configure a **Captive Portal Detection Message**.



These features require Content Release version 607-3486 or later.

STEP 4 | Specify additional GlobalProtect connection settings.



With single sign-on (SSO) enabled (the default), the GlobalProtect agent uses the user's Windows login credentials to automatically authenticate to and connect to the GlobalProtect portal and gateway. GlobalProtect with SSO enabled also allows for the GlobalProtect agent to wrap third-party credentials to ensure that Windows users can authenticate and connect, even when a third-party credential provider is being used to wrap the Windows login credentials.

In the App Configurations area, configure any of the following options:

- (**Windows only**) Set **Use Single Sign-On** to **No** to disallow GlobalProtect to use the Windows login credentials to automatically authenticate the user upon login to Active Directory.
- Enter the **Maximum Internal Gateway Connection Attempts** to specify the number of times the GlobalProtect agent should retry the connection to an internal gateway after the first attempt fails (range is 0-100; 4 or 5 is recommended; default is 0, which means the GlobalProtect agent does not retry the connection). By increasing the value, you enable the agent to connect to an internal gateway that is temporarily down or unreachable during the first connection attempt but comes back up before the specified number of retries are exhausted. Increasing the value also ensures that the internal gateway receives the most up-to-date user and host information.
- Enter the **GlobalProtect App Config Refresh Interval (hours)** to specify the number of hours the GlobalProtect portal waits before it initiates the next refresh of a client's configuration (range is 1-168; default is 24).
- Specify whether to **Retain Connection on Smart Card Removal**. By default, the option is set to **Yes**, meaning GlobalProtect retains the tunnel when a user removes a smart card containing a client certificate. To terminate the tunnel, set this option to **No**. The decision on whether to retain the connection depends on your security requirements.



This feature requires Content Release version 590-3397 or a later version.

- Configure an **Automatic Restoration of VPN Connection Timeout** to specify the action GlobalProtect takes when the tunnel is disconnected by entering a timeout value in minutes from 0 to 180; default is 30. A value of 0 disables this feature so that GlobalProtect does not attempt to reconnect after the tunnel is disconnected. When you specify a value of 1-180 minutes, GlobalProtect attempts to reestablish the tunnel connection if the tunnel is down for a period of time which does not exceed the timeout value you specify here. For example, with a timeout value of 30 minutes, GlobalProtect does not attempt to reconnect if the tunnel is disconnected for 45 minutes. However, if the tunnel is

© 2017 Palo Alto Networks, Inc.

disconnected for 15 minutes, GlobalProtect attempts to reconnect because the number of minutes has not exceeded the timeout value.

When you enable **Automatic Restoration of VPN Connection Timeout**, you can also adjust the amount of time in seconds GlobalProtect waits between attempts to restore the connection by configuring the **Wait Time Between VPN Connection Restore Attempts**. Range is 1 to 60 seconds; the default is 5.

With Always-On VPN, if a user switches from an external network to an internal network before the timeout value expires, GlobalProtect does not perform network discovery. As a result, GlobalProtect restores the connection to the last known external gateway. To trigger internal host detection, the user must select Rediscover Network from the GlobalProtect console.

STEP 5 | Configure the menus and UI views that are available to users who have this agent configuration.

Configure any or all of the following options:

- If you want users to be able to see only basic status information within the application, set **Enable Advanced View** to **No**. By default, the advanced view is enabled. It allows users to see detailed statistical, host, and troubleshooting information and to perform certain tasks, such as changing their password.
- If you want hide the GlobalProtect agent on end-user systems, set **Display GlobalProtect Icon** to **No**. When the icon is hidden, users cannot perform other tasks such as changing passwords, rediscovering the network, resubmitting host information, viewing troubleshooting information, or performing an on-demand connection. However, HIP notification messages, login prompts, and certificate dialogs will still display as necessary for interacting with the end user.
- To prevent users from performing a network rediscovery, set the **Enable Rediscover Network Option** to **No**. When you disable the option, it is grayed out in the GlobalProtect menu.
- To prevent users from manually resubmitting HIP data to the gateway, set **Enable Resubmit Host Profile Option** to **No**. This option is enabled by default, and is useful in cases where HIP-based security policy prevents users from accessing resources because it allows the user to fix the compliance issue on the computer and then resubmit the HIP.
- (**Windows only**) To allow GlobalProtect to display notifications in the notification area (system tray), set **Show System Tray Notifications** to **Yes**.
- To create a custom message to display to users when their password is about to expire configure the **Custom Password Expiration Message (LDAP Authentication Only)**. The maximum message length is 200 characters.

STEP 6 | Define what the end users with this configuration can do in their client.

- Set **Allow User to Change Portal Address** to **No** to disable the **Portal** field on the **Home** tab in the GlobalProtect agent. Because the user will then be unable to specify a portal to which to connect, you must supply the default portal address in the Windows registry (`HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\PanSetup` with key `Portal`) or the Mac plist (`/Library/Preferences/com.paloaltonetworks.GlobalProtect.settings.plist` with key `Portal` under dictionary `PanSetup`). For more information, see Deploy Agent Settings Transparently.
- To prevent users from dismissing the welcome page, set **Allow User to Dismiss Welcome Page** to **No**. Otherwise, when set to **Yes**, the user can dismiss the welcome page and prevent GlobalProtect from displaying the page after subsequent logins.

STEP 7 | Specify whether users can disable the GlobalProtect agent.

© 2017 Palo Alto Networks, Inc.

The **Allow User to Disable GlobalProtect** option applies to agent configurations that have the **Connect Method** set to **User-Logon (Always On)**. In user-logon mode, the agent or app automatically connects to GlobalProtect as soon as the user logs in to the endpoint. This mode is sometimes referred to as "always on," which is why the user must override this behavior to disable GlobalProtect client.

By default, this option is set to **Allow** which permits users to disable GlobalProtect without providing a comment, passcode, or ticket number.

If the agent icon is not visible, users are not able to disable the GlobalProtect client. See 5 for details.

- To prevent users with the user-logon connect method from disabling GlobalProtect, set **Allow User to Disable GlobalProtect** to **Disallow**.
- To allow users to disable GlobalProtect if they provide a passcode, set **Allow User to Disable GlobalProtect** to **Allow with Passcode**. Then, in the Disable GlobalProtect App area, enter (and confirm) the **Passcode** that the end users must supply.
- To allow users to disconnect if they provide a ticket, set **Allow User to Disable GlobalProtect** to **Allow with Ticket**. With this option, the disconnect action triggers the agent to generate a Request Number. The end user must then communicate the Request Number to the administrator. The administrator then clicks **Generate Ticket** on the **Network** > **GlobalProtect** > **Portals** page and enters the request number from the user to generate the ticket. The administrator then provides the ticket to the end user, who enters it into the Disable GlobalProtect dialog to enable the agent to disconnect.



- To limit the number of times users can disable the GlobalProtect client, enter a value in the **Max Times User Can Disable** field in the Disable GlobalProtect App area. A value of 0 (the default) indicates that users are not limited in the number of times they can disable the client.
- To restrict how long the user may be disconnected, enter a value (in minutes) in the **User Can Disable Timeout (min)** field in the Disable GlobalProtect App area. A value of 0 (the default) means that there is no restriction on how long the user can keep the client disabled.

STEP 8 | Configure the certificate settings and behavior for the users that receive this configuration.

- **Client Certificate Store Lookup**—Select which store the agent should use to look up client certificates. **User** certificates are stored in the Current User certificate store on Windows and in the Personal Keychain on Mac OS. **Machine** certificates are stored in the Local Computer certificate store on Windows and in the System Keychain on Mac OS. By default, the agent looks for **User and machine** certificates in both places.
- **SCEP Certificate Renewal Period (days)**—With SCEP, the portal can request a new client certificate before the certificate expires. This time before the certificate expires is the optional *SCEP certificate renewal period*. During a configurable number of days before a client certificate expires, the portal can request a new certificate from the SCEP server in your enterprise PKI (range is 0-30; default is 7). A value of 0 means the portal does not automatically renew the client certificate when it refreshes the agent configuration.

  For an agent or app to obtain the new certificate during the renewal period, the user must log in to the GlobalProtect client. For example, if a client certificate has a lifespan of 90 days, the certificate renewal period is 7 days, and the user logs in during the final 7 days of the certificate lifespan, the

© 2017 Palo Alto Networks, Inc.

portal acquires a new certificate and deploys it along with a fresh agent configuration. For more information, see Deploy User-Specific Client Certificates for Authentication.

- **Extended Key Usage OID for Client Certificate** (Windows and Mac endpoints only)—Use this option only if you enabled client authentication, expect multiple client certificates to be present on the endpoint, and have identified a secondary purpose by which you can filter the client certificates. This option enables you to specify a secondary purpose for a client certificate using the associated object identifier (OID). For example, to display only client certificates which also have a purpose of Server Authentication, enter the OID 1.3.6.1.5.5.7.3.1. When the GlobalProtect agent finds only one client certificate which matches the secondary purpose, GlobalProtect automatically selects and authenticates using that certificate. Otherwise, GlobalProtect prompts the user to select the client certificate from the filtered list of client certificates which match the criteria. For more information including a list of common certificate purposes and OIDs, see the PAN-OS 7.1 New Feature'sGuide.
- If you do not want the agent to establish a connection with the portal when the portal certificate is not valid, set **Allow User to Continue with Invalid Portal Server Certificate** to **No**. Keep in mind that the portal provides the agent configuration only; it does not provide network access and therefore security to the portal is less critical than security to the gateway. However, if you have deployed a trusted server certificate for the portal, deselecting this option can help prevent man-in-the-middle (MITM) attacks.

STEP 9 | Specify whether users receive login prompts when multi-factor authentication is required to access sensitive network resources.

For internal gateway connections, sensitive network resources (for example, financial applications or software development applications) may require additional authentication. You can configure GlobalProtect clients to display the authentication prompts required to access these resources. Refer to Configure GlobalProtect to Facilitate Multi-Factor Authentication Notifications for more information.

- Set **Enable Inbound Authentication Prompts from MFA Gateways** to **Yes**. To support multi-factor authentication (MFA), a GlobalProtect client must receive and acknowledge UDP prompts that are inbound from the gateway. Select **Yes** to enable a GlobalProtect client to receive and acknowledge the prompt. By default, the value is set to **No** meaning GlobalProtect will block UDP prompts from the gateway.
- In **Network Port for Inbound Authentication Prompts (UDP)**, specify the port number a GlobalProtect client uses to receive inbound authentication prompts from MFA gateways. The default port is 4501. To change the port, specify a number from 1 to 65535.
- In **Trusted MFA Gateways**, specify the list of authentication gateways a GlobalProtect client will trust for multi-factor authentication. When a GlobalProtect client receives a UDP message on the specified network port, GlobalProtect displays an authentication message only if the UDP prompt comes from a trusted gateway.
- Configure the **Inbound Authentication Message** (for example: `You have attempted to access a protected resource that requires additional authentication. Proceed to authenticate at)`. When users try to access a resource that requires additional authentication, GlobalProtect receives an inbound authentication prompt and displays this message. GlobalProtect automatically appends the URL for the Authentication Portal page you specify when you configure multi-factor authentication to the inbound authentication message.

STEP 10 | (**Windows only**) Configure settings for Windows-based endpoints that receive this configuration.

- **Resolve All FQDNs Using DNS Servers Assigned by the Tunnel (Windows Only)**— Configure the DNS resolution preferences for the GlobalProtect tunnel. Select **No** to allow Windows endpoints to send DNS queries to the DNS server set on the physical adapter if the initial query to the DNS server configured on the gateway is not resolved. This option retains the native Windows behavior to query all DNS servers on all adapters recursively but can result in long wait times to resolve some DNS queries. Select **Yes** (default) to allow Windows endpoints to resolve all DNS queries with the DNS

© 2017 Palo Alto Networks, Inc.

servers you configure on the gateway instead of allowing the endpoint to send some DNS queries to the DNS servers set on the physical adapter.

This features requires Content Release version 731 or later and is available for GlobalProtect agent 4.0.3 and later.

To configure DNS settings for GlobalProtect agent 4.0.2 and earlier releases, use the **Update DNS Settings at Connect** option.

- **Update DNS Settings at Connect**—Select **Yes** to enable the Windows endpoint to resolve all DNS queries with the DNS servers you configure for the gateway instead of the DNS servers set for the physical adapter on the endpoint. When you enable this option, GlobalProtect strictly enforces the gateway DNS settings and overrides the static settings for all physical adapters. This is useful when a Windows endpoint fails to resolve a DNS query sent to the DNS server configured on the physical adapter instead of on the GlobalProtect tunnel adapter. Select **No** (the default) to allow Windows endpoints to send DNS queries to the DNS server set on the physical adapter if the initial query to the DNS server configured on the gateway is not resolved. This option retains the native Windows behavior to query all DNS servers on all adapters recursively but can result in long wait times to resolve some DNS queries.

  This feature is deprecated in 4.0.3 and later releases. To configure DNS resolution settings for GlobalProtect agent 4.0.3 and later releases, use the Resolve All FQDNs Using DNS Servers Assigned by the Tunnel (Windows Only) option.
- **Send HIP Report Immediately if Windows Security Center (WSC) State Changes**—Select **No** to prevent the GlobalProtect agent from sending HIP data when the status of the Windows Security Center (WSC) changes. Select **Yes** (default) to immediately send HIP data when the status of the WSC changes.
- **Detect Proxy for Each Connection**—Select **No** to auto-detect the proxy for the portal connection and use that proxy for subsequent connections. Select **Yes** (default) to auto-detect the proxy at every connection.
- **Clear Single Sign-On Credentials on Logout**—Select **No** to keep single sign-on credentials when the user logs out. Select **Yes** (default) to clear them and force the user to enter credentials upon the next login.
- **Use Default Authentication on Kerberos Authentication Failure**—Select **No** to use only Kerberos authentication. Select **Yes** (default) to retry using the default authentication method after authentication using Kerberos fails.

STEP 11 | If your endpoints frequently experience latency or slowness when connecting to the GlobalProtect portal or gateways, consider adjusting the portal and TCP timeout values.

To allow more time for your endpoints to connect to or receive data from the portal or gateway, increase the timeout values, as needed. Keep in mind that increasing the values can result in longer wait times if the GlobalProtect agent is unable to establish the connection. In contrast, decreasing the values can prevent the GlobalProtect agent from establishing a connection when the portal or gateway does not respond before the timeout expires.

Configure values for any of the following options:

- **Portal Connection Timeout (sec)**—The number of seconds (between 1 and 600) before a connection request to the portal times out due to no response from the portal. When your firewall is running Applications and Threats content versions earlier than 777-4484, the default is 30. Starting with content version 777-4484, the default is 5.
- **TCP Connection Timeout (sec)**—The number of seconds (between 1 and 600) before a TCP connection request times out due to unresponsiveness from either end of the connection. When your firewall is running Applications and Threats content versions earlier than 777-4484, the default is 60. Starting with content version 777-4484, the default is 5.

© 2017 Palo Alto Networks, Inc.

- **TCP Receive Timeout (sec)**—The number of seconds before a TCP connection times out due to the absence of some partial response of a TCP request (range is 1-600; default is 30).

STEP 12 | Specify whether remote desktop connections are permitted over existing VPN tunnels by specifying the **User Switch Tunnel Rename Timeout**. When a new user connects to a Windows machine using Remote Desktop Protocol (RDP), the gateway reassigns the VPN tunnel to the new user. The gateway can then enforce security policies on the new user.

Allowing remote desktop connections over VPN tunnels can be useful in situations where an IT administrator needs to access a remote end-user system using RDP.

By default, the **User Switch Tunnel Rename Timeout** field is set to 0 meaning the GlobalProtect gateway terminates the connection if a new user authenticates over the VPN tunnel. To modify this behavior, configure a timeout value from 1 to 600 seconds. If the new user does not log in to the gateway before the timeout value expires, the GlobalProtect gateway terminates the VPN tunnel assigned to the first user.



Changing the User Switch Tunnel Rename Timeout value only affects the RDP tunnel and does not rename a pre-logon tunnel when configured.

STEP 13 | Specify how GlobalProtect agent upgrades occur.

If you want to control when users can upgrade, you can customize the agent upgrade on a per-configuration basis. For example, if you want to test a release on a small group of users before deploying it to your entire user base, you can create a configuration that applies to users in your IT group only, thus allowing them to upgrade and test and disable upgrade in all other user/group configurations. Then, after you have thoroughly tested the new version, you can modify the agent configurations for the rest of your users to allow the upgrade.

By default, the **Allow User to Upgrade GlobalProtect App** field is set to **prompt** the end user to upgrade. To modify this behavior, select one of the following options:

- **Allow Transparently**—Upgrades occur automatically without interaction with the user. Upgrades can occur when the user is working remotely or connected from within the corporate network.
- **Internal**—Upgrades occur automatically without interaction with the user, provided the user is connected from within the corporate network. This setting is recommended to prevent slow upgrades in low-bandwidth situations. When a user connects outside the corporate network, the upgrade is postponed and re-activated later when the user connects from within the corporate network. You must configure internal gateways and internal host detection to use this option.
- To prevent agent upgrades, select **Disallow**.
- To allow end users to initiate agent upgrades, select **Allow Manually**. In this case, the user would select the **Check Version** option in the agent to determine if there is a new agent version and then upgrade if desired. Note that this option will not work if the GlobalProtect agent is hidden from the user. See 5 for details on the **Display GlobalProtect Icon** option.



Upgrades for Allow Transparently and Internal occur only if the GlobalProtect software version on the portal is more recent than the GlobalProtect software version on the endpoint. For example, a GlobalProtect 3.1.3 agent connecting to a GlobalProtect 3.1.1 portal is not upgraded.

STEP 14 | Specify whether to display a welcome page upon successful login.

A welcome page can be a useful way to direct users to internal resources that they can only access when connected to GlobalProtect, such as your Intranet or other internal servers.

By default, the only indication that the agent has successfully connected to GlobalProtect is a balloon message that displays in the system tray/menubar.

© 2017 Palo Alto Networks, Inc.

To display a welcome page after a successful login select **factory-default** from the **Welcome Page** drop-down on the right. GlobalProtect displays the welcome page in the default browser on Windows, Mac, and Chromebook endpoints, or within the GlobalProtect app on mobile devices. You can also select a custom welcome page that provides information specific to your users, or to a specific group of users (based on which portal configuration gets deployed). For details on creating custom pages, see Customize the GlobalProtect Portal Login, Welcome, and Help Pages.

STEP 15 | Save the agent configuration settings.

1. If you are done creating agent configurations, click **OK** to close the Configs dialog. Otherwise, for instructions on completing the agent configurations, return to Define the GlobalProtect Agent Configurations.
2. If you are done configuring the portal, click **OK** to close the GlobalProtect Portal Configuration dialog.
3. When you finish the portal configuration, **Commit** the changes.

Customize the GlobalProtect Portal Login, Welcome, and Help Pages

GlobalProtect provides default login, welcome, and/or help pages. However, you can create your own custom pages with your corporate branding, acceptable use policies, and links to your internal resources.

You can alternatively disable browser access to the portal login page in order to prevent unauthorized attempts to authenticate to the GlobalProtect portal (configure the Portal Login Page > Disable option from Network > GlobalProtect > Portals > <portal_config > General). With the portal login page disabled, you can instead use a software distribution tool, such as Microsoft's System Center Configuration Manager (SCCM), to allow your users to download and install the GlobalProtect agent.

STEP 1 | Export the default portal login, welcome, or help page.

1. Select **Device** > **Response Pages**.
2. Select the link for the type of GlobalProtect portal page.
3. Select the **Default** predefined page and click **Export**.

STEP 2 | Edit the exported page.

1. Use the HTML text editor of your choice to edit the page.
2. If you want to edit the logo image that is displayed, host the new logo image on a web server that is accessible from the remote GlobalProtect clients. For example, edit the following line in the HTML to point to the new logo image:

   ```
   <img src="http://cdn.slidesharecdn.com/ Acme-logo-96x96.jpg?1382722588"/>
   ```
3. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding.

STEP 3 | Import the new page(s).

1. Select **Device** > **Response Pages**.
2. Select the link for the type of GlobalProtect portal page.
3. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.
4. (**Optional**) Select the virtual system on which this page will be used from the **Destination** drop-down or select shared (default) to make it available to all virtual systems.
5. Click **OK** to import the file.

STEP 4 | Configure the portal to use the new page(s).

© 2017 Palo Alto Networks, Inc.

- **Portal Login Page** and **App Help Page**:

1. Select **Network** > **GlobalProtect** > **Portals** and select the portal to which you want to add the login page.
2. On the **General** tab, select the new page from the relevant drop-down in the Appearance area.

- **Custom Welcome Page**:

3. Select **Network** > **GlobalProtect** > **Portals** and select the portal to which you want to add the login page.
4. On the **Agent** tab, select the agent configuration to which you want to add the welcome page.
5. Select the **App** tab, and select the new page from the **Welcome Page** drop-down.
6. Click **OK** to save the agent configuration.

STEP 5 | Save the portal configuration.

Click **OK** and then **Commit** your changes.

STEP 6 | Verify that the new page displays.

- **Test the login page**—Open a browser, go to the URL for your portal (be sure you do not add the :4443 port number to the end of the URL or you will be directed to the web interface for the firewall). For example, enter **`https://myportal`** rather than **`https://myportal:4443`**.

The new portal login page will display.



- **Test the help page**—Right-click the GlobalProtect icon in the notification area (system tray), and select **Help**. The new help page will display.
- **Test the welcome page**—Right-click the GlobalProtect icon in the notification area (system tray), and select **Welcome Page**. The new welcome page will display.

GlobalProtect Clients

- Deploy the GlobalProtect Client Software on page 103
- Define the GlobalProtect Agent Configurations on page 86
- Customize the GlobalProtect Agent on page 91
- Deploy Agent Settings Transparently on page 115
- GlobalProtect Clientless VPN on page 135

© 2017 Palo Alto Networks, Inc.

Deploy the GlobalProtect Client Software

In order to connect to GlobalProtect™, an end host must be running GlobalProtect client software. The software deployment method depends on the type of client as follows:

- **Mac OS and Microsoft Windows endpoints**—Require the GlobalProtect agent software, which is distributed by the GlobalProtect portal. To enable the software for distribution, you must download the version you want the hosts in your network to use to the firewall hosting your GlobalProtect portal and then activate the software for download. For instructions on how to download and activate the agent software on the firewall, see Deploy the GlobalProtect Agent Software.
- **Windows 10 phone and Windows 10 UWP endpoints**—Require the GlobalProtect app. As with other mobile device apps, the end user must download the GlobalProtect app from the Microsoft Store. For instructions on how to download and test the GlobalProtect app installation, see Download and Install the GlobalProtect Mobile App.
- **iOS and Android endpoints**—Require the GlobalProtect app. As with other mobile device apps, the end user must download the GlobalProtect app either from the Apple AppStore (iOS devices) or from Google Play (Android devices). For instructions on how to download and test the GlobalProtect app installation, see Download and Install the GlobalProtect Mobile App.
- **Chromebooks**—Require the GlobalProtect app for Chrome OS. Similar to the download process for mobile device apps, the end user can download the GlobalProtect app from the Chrome Web Store. You can also deploy the app to a managed Chromebook using the Chromebook Management Console. For instructions on how to download and test the GlobalProtect app installation, Download and Install the GlobalProtect App for Chrome OS.

For more details, see What Client OS Versions are Supported with GlobalProtect?



As an alternative to deploying GlobalProtect client software, you can configure the GlobalProtect portal to provide secure remote access to common enterprise web applications that use HTML, HTML5, and Javascript technologies. Users have the advantage of secure access from SSL-enabled web browsers without installing GlobalProtect client software. Refer to GlobalProtect Clientless VPN.

Deploy the GlobalProtect Agent Software

There are several ways to deploy the GlobalProtect agent software:

- **Directly from the portal**—Download the agent software to the firewall hosting the portal and activate it so that end users can install the updates when they connect to the portal. This option provides flexibility in that it allows you to control how and when end users receive updates based on the agent configuration settings you define for each user, group, and/or operating system. However, if you have a large number of agents that require updates, it could put extra load on your portal. See Host Agent Updates on the Portal for instructions.
- **From a web server**—If you have a large number of hosts that will need to upgrade the agent simultaneously, consider hosting the agent updates on a web server to reduce the load on the firewall. See Host Agent Updates on a Web Server for instructions.
- **Transparently from the command line**—For Windows clients, you can automatically deploy agent settings in the Windows Installer (Msiexec). However, to upgrade to a later agent version using Msiexec, you must first uninstall the existing agent. In addition, Msiexec allows for deployment of agent settings directly on the endpoints by setting values in the Windows registry or Mac plist. See Deploy Agent Settings Transparently.
- **Using group policy rules**—In Active Directory environments, the GlobalProtect Agent can also be distributed to end users, using active directory group policy. AD Group policies allow modification of Windows host computer settings and software automatically. Refer to the article at http://

© 2017 Palo Alto Networks, Inc.

support.microsoft.com/kb/816102 for more information on how to use Group Policy to automatically distribute programs to host computers or users.

- **From a mobile endpoint management system**—If you use an mobile management system such as an MDM or EMM to manage your mobile devices, you can use the system to deploy and configure the GlobalProtect app. See Mobile Endpoint Management.

Host Agent Updates on the Portal

The simplest way to deploy the GlobalProtect agent software is to download the new agent installation package to the firewall that is hosting your portal and then activate the software for download to the agents connecting to the portal. To do this automatically, the firewall must have a service route that enables it to access the Palo Alto Networks Update Server. If the firewall does not have access to the Internet, you can manually download the agent software package from the Palo Alto Networks Software Updates support site using an Internet-connected computer and then manually upload it to the firewall.



You must have a valid Palo Alto Networks account to log in to and download software from the Software Updates page. If you cannot log in and need assistance, go to https://www.paloaltonetworks.com/support/tabs/overview.html.)

You define how the agent software updates are deployed in the agent configurations you define on the portal—whether they happen automatically when the agent connects to the portal, whether the user is prompted to upgrade the agent, or whether the end user can manually check for and download a new agent version. For details on creating an agent configuration, see Define the GlobalProtect Agent Configurations on page 86.

STEP 1 | Launch the web interface on the firewall hosting the GlobalProtect portal and go to the GlobalProtect Client page.

Select **Device** > **GlobalProtect Client**.

STEP 2 | Check for new agent software images.

- If the firewall has access to the Update Server, click **Check Now** to check for the latest updates. If the value in the **Action** column is **Download** it indicates that an update is available.
- If the firewall does not have access to the Update Server, go to the Palo Alto Networks Software Updates support site and **Download** the file to your computer. Then go back to the firewall to manually **Upload** the file.



You must have a valid Palo Alto Networks account to log in to and download software from the Software Updates page. If you cannot log in and need assistance, go to: https://www.paloaltonetworks.com/support/tabs/overview.html)

STEP 3 | Download the agent software image.



If your firewall does not have Internet access from the management port, you can download the agent update from the Palo Alto Networks Support Site: (https://www.paloaltonetworks.com/support/tabs/overview.html).

You can then manually **Upload** the update to your firewall and then activate **Activate From File**.

Locate the agent version you want and then click **Download**. When the download completes, the value in the **Action** column changes to **Activate**.

© 2017 Palo Alto Networks, Inc.

If you manually uploaded the agent software as detailed in 2 on page 104, the Action column will not update. Continue to the next step for instructions on activating an image that was manually uploaded.

STEP 4 | Activate the agent software image so that end users can download it from the portal.

Only one version of agent software image can be activated at a time. If you activate a new version, but have some agents that require a previously activated version, you will have to activate the required version again to enable it for download.

- If you downloaded the image automatically from the Update Server, click **Activate**.
- If you manually uploaded the image to the firewall, click **Activate From File** and then select the **GlobalProtect Client File** you uploaded from the drop-down. Click **OK** to activate the selected image. You may need to refresh the screen before the version displays as **Currently Activated**.

Host Agent Updates on a Web Server

If you have a large number of endpoints that will need to install and/or update the GlobalProtect agent software, consider hosting the GlobalProtect agent software images on an external web server. This helps reduce the load on the firewall when users connect to download the agent. To use this feature, the firewall hosting the portal must be running PAN-OS 4.1.7 or a later release.

STEP 1 | Download the version of the GlobalProtect agent that you plan to host on the web server to the firewall and activate it.

Follow the steps for downloading and activating the agent software on the firewall as described in Host Agent Updates on the Portal.

STEP 2 | Download the GlobalProtect agent image you want to host on your web server.

You should download the same image that you activated on the portal.

From a browser, go to the Palo Alto Networks Software Updates site and **Download** the file to your computer.

STEP 3 | Publish the files to your web server.

Upload the image file(s) to your web server.

STEP 4 | Redirect the end users to the web server.

On the firewall hosting the portal, log in to the CLI and enter the following operational mode commands:

```
> set global-protect redirect on
> set global-protect redirect location <path>
```

where *<path>* is the path is the URL to the folder hosting the image, for example **https://acme/GP**.

STEP 5 | Test the redirect.

1. Launch your web browser and go to the following URL:

```
https://<portal address or name>
```

For example, **https://gp.acme.com**.

© 2017 Palo Alto Networks, Inc.

2. On the portal login page, enter your user **Name** and **Password** and then click **Login**. After successful login, the portal should redirect you to the download.

Test the Agent Installation

Use the following procedure to test the agent installation.

STEP 1 | Create an agent configuration for testing the agent installation.

When initially installing the GlobalProtect agent software on the endpoint, the end user must be logged in to the system using an account that has administrative privileges. Subsequent agent software updates do not require administrative privileges.

As a best practice, create an agent configuration that is limited to a small group of users, such as administrators in the IT department responsible for administering the firewall:

1. Select **Network** > **GlobalProtect** > **Portals** and select the portal configuration to edit.
2. Select the **Agent** tab and either select an existing configuration or **Add** a new configuration to deploy to the test users/group.
3. On the **User/User Group** tab, click **Add** in the User/User Group section, select the user or group who will be testing the agent, and then click **OK**.
4. On the **Agent** tab, make sure **Agent Upgrade** is set to **prompt** and then click **OK** to save the configuration.
5. (**Optional**) Select the agent configuration you just created/modified and click **Move Up** so that it is before any more generic configurations you have created.
6. **Commit** the changes.

STEP 2 | Log in to the GlobalProtect portal.

1. Launch your web browser and go to the following URL:

```
https://<portal address or name>
```

For example, **`https://gp.acme.com`**.

2. On the portal login page, enter your user **Name** and **Password** and then click **Login**.



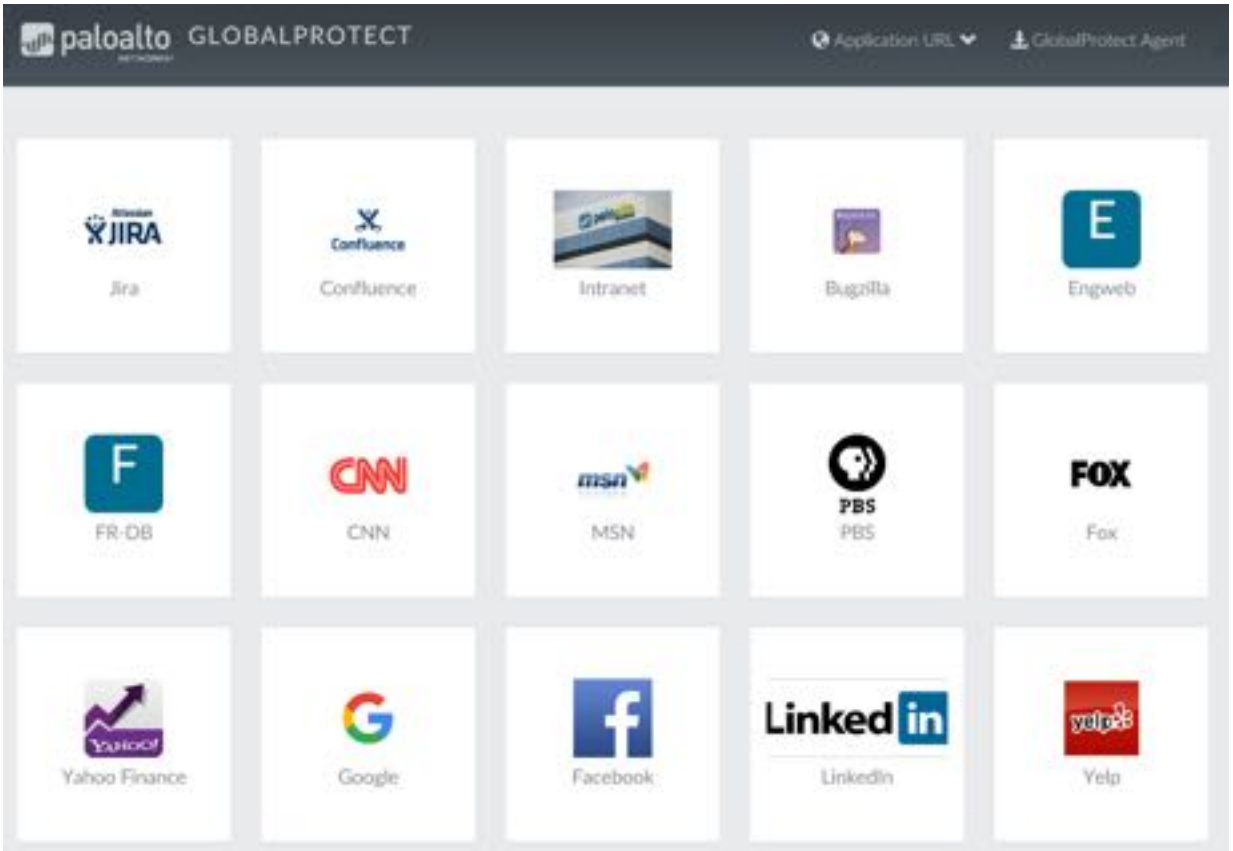
STEP 3 | Navigate to the agent download page.

In most cases, you will see an agent download page when you log in to the portal. Use this page to download the latest agent software package.

© 2017 Palo Alto Networks, Inc.



If you have enabled GlobalProtect Clientless VPN access, you will see an applications page (instead of the agent download page) when you log in to the portal. Select **GlobalProtect Agent** to open the download page.



STEP 4 | Download the agent.

1. Click the link that corresponds to the operating system you are running on your computer to begin the download.

© 2017 Palo Alto Networks, Inc.



2. When prompted to run or save the software, click **Run**.
3. When prompted, click **Run** to launch the GlobalProtect Setup Wizard.

When initially installing the GlobalProtect agent software on the endpoint, the end user must be logged in to the system using an account that has administrative privileges. Subsequent agent software updates do not require administrative privileges.

STEP 5 | Complete the GlobalProtect agent setup.

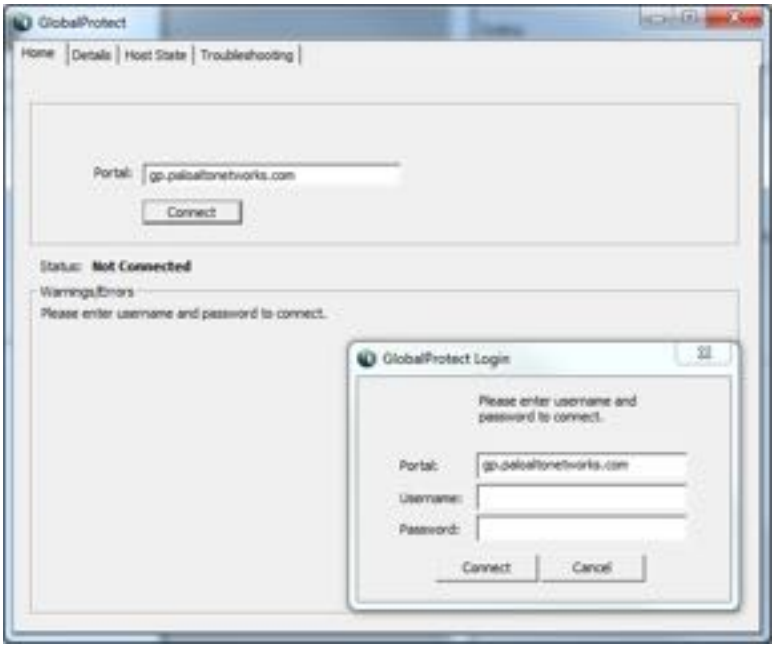
1. From the GlobalProtect Setup Wizard, click **Next**.
2. Click **Next** to accept the default installation folder

   (`C:\Program Files\Palo Alto Networks\GlobalProtect`) or **Browse** to choose a new location and then click **Next** twice.
3. After the installation successfully completes, click **Close**. The GlobalProtect agent will automatically start.

STEP 6 | Log in to GlobalProtect.

Enter the FQDN or IP address of the **Portal** and **Connect**. If prompted, enter your **Username** and **Password** and **Connect**. If authentication is successful, the agent will connect to GlobalProtect. Use the agent to access resources on the corporate network as well as external resources, as defined in the corresponding security polices.

© 2017 Palo Alto Networks, Inc.



To deploy the agent to end users, create agent configurations for the user groups for which you want to enable access and set the **Agent Upgrade** settings appropriately and then communicate the portal address. See Define the GlobalProtect Agent Configurations for details on setting up agent configurations.

Download and Install the GlobalProtect Mobile App

The GlobalProtect app provides a simple way to extend the enterprise security policies out to mobile devices. As with other remote hosts running the GlobalProtect agent, the mobile app provides secure access to your corporate network over an IPSec or SSL VPN tunnel. The app will automatically connect to the gateway that is closest to the end user's current location. In addition, traffic to and from the mobile device is automatically subject to the same security policy enforcement as other hosts on your corporate network. Like the GlobalProtect agent, the app collects information about the host configuration and can use this information for enhanced HIP-based security policy enforcement.

There are two primary methods for installing the GlobalProtect app: You can deploy the app from your third-party MDM and transparently push the app to your managed devices; or, you can install the app directly from the official store for your device:

- iOS endpoints—App Store
- Android endpoints—Google Play
- Windows 10 phones and Windows 10 UWP endpoints—Microsoft Store
- Chromebooks—For details on installing the GlobalProtect app for Chrome OS, see Download and Install the GlobalProtect App for Chrome OS.

This workflow describes how to install the GlobalProtect app directly on the mobile device. For instructions on how to deploy the GlobalProtect app from AirWatch, see Deploy the GlobalProtect Mobile App Using AirWatch.

STEP 1 | Create an agent configuration for testing the app installation.

As a best practice, create an agent configuration that is limited to a small group of users, such as administrators in the IT department responsible for administering the firewall:

© 2017 Palo Alto Networks, Inc.

1. Select **Network** > **GlobalProtect** > **Portals** and select the portal configuration to edit.
2. Select the **Agent** tab and either select an existing configuration or **Add** a new configuration to deploy to the test users/group.
3. On the **User/User Group** tab, click **Add** in the User/User Group section and then select the user or group who will be testing the agent.
4. In the OS section, select the app you are testing (**iOS**, **Android**, or **WindowsUWP**).
5. (**Optional**) Select the agent configuration you just created/modified and click **Move Up** so that it is before any more generic configurations you have created.
6. **Commit** the changes.

STEP 2 | From the mobile device, follow the prompts to download and install the app.

- On Android devices, search for the app on Google Play.
- On iOS devices, search for the app at the App Store.
- On Windows 10 UWP devices, search for the app at the Microsoft Store.

STEP 3 | Launch the app.

When successfully installed, the GlobalProtect app icon displays on the device's Home screen. To launch the app, tap the icon. When prompted to enable GlobalProtect VPN functionality, tap **OK**.



STEP 4 | Connect to the portal.

1. When prompted, enter the **Portal** name or address, **Username**, and **Password**. The portal name must be an FQDN and it should not include the https:// at the beginning.

© 2017 Palo Alto Networks, Inc.



2. Tap **Connect** and verify that the app successfully establishes a VPN connection to GlobalProtect.

   If a third-party mobile endpoint management system is configured, the app will prompt you to enroll.

Download and Install the GlobalProtect App for Chrome OS

The GlobalProtect app for Chrome OS provides a simple way to extend the enterprise security policies out to Chromebooks. As with other remote hosts running the GlobalProtect agent, the app provides secure access to your corporate network over an IPSec or SSL VPN tunnel. After the user initiates a connection, the app will connect to the gateway that is closest to the end user's current location. In addition, traffic to and from the Chromebook is automatically subject to the same security policy enforcement as other hosts on your corporate network. Like the GlobalProtect agent, the app collects information about the host configuration and can use this information for enhanced HIP-based security policy enforcement.

Use the following procedures to install and test the GlobalProtect app for Chrome OS.

- Install the GlobalProtect App from the Chrome Web Store on page 111
- Deploy the GlobalProtect App Using the Chromebook Management Console on page 112
- Test the GlobalProtect app for Chrome OS on page 113

Install the GlobalProtect App from the Chrome Web Store

You can install the GlobalProtect app on a Chromebook by downloading the app from the Chrome Web Store. As an alternative you can Deploy the GlobalProtect App Using the Chromebook Management Console on page 112.

STEP 1 | Create an agent configuration for testing the app installation.



As a best practice, create an agent configuration that is limited to a small group of users, such as administrators in the IT department and who responsible for administering the firewall.

1. Select **Network** > **GlobalProtect** > **Portals** and select the portal configuration to edit.
2. Select the **Agent** tab and either select an existing configuration or **Add** a new configuration to deploy to the test users/group.
3. On the **User/User Group** tab, click **Add** in the User/User Group section and then select the user or group that will test the agent.
4. In the OS area, select the app you are testing (Chrome) and click **OK**.
5. (**Optional**) Select the agent configuration you just created or modified and click **Move Up** so that it is before any more generic configurations you have created.
6. **Commit** the changes.

© 2017 Palo Alto Networks, Inc.

STEP 2 | Install the GlobalProtect app for Chrome OS.

You can also force-install the app on managed Chromebooks using the Chromebook Management Console. See Deploy the GlobalProtect App Using the Chromebook Management Console on page 112.

1. From the Chromebook, search for the app in the Chrome Web Store or go directly to the GlobalProtect app page.
2. Click **Add to Chrome** and then follow the prompts to download and install the app.

STEP 3 | Launch the app.

When successfully installed, the Chrome App Launcher displays the GlobalProtect app icon in the list of apps**.** To launch the app, click the icon.

STEP 4 | Configure the portal.

1. When prompted, enter the IP address or FQDN of the **Portal**. The portal should not include the https:// at the beginning.
2. Click **Add Connection** to add the GlobalProtect VPN configuration.

   The app displays the home screen after it adds the VPN configuration to the Internet connection settings of your Chromebook but does not initiate a connection.

STEP 5 | Test the connection.

Test the GlobalProtect app for Chrome OS on page 113

Deploy the GlobalProtect App Using the Chromebook Management Console

The Chromebook Management Console enables you to manage Chromebook settings and apps from a central, web-based location. From the console, you can deploy the GlobalProtect app to Chromebooks and customize VPN settings.

Use the following workflow to manage policies and settings for the GlobalProtect app for Chrome OS:

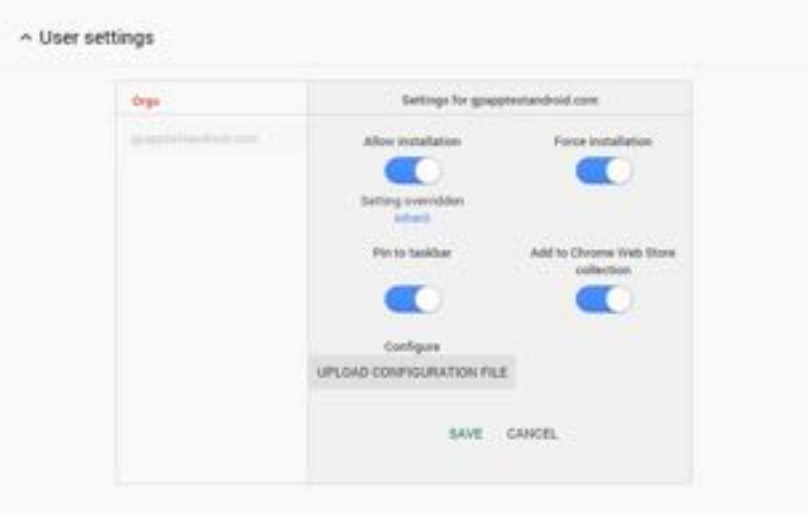
STEP 1 | View the user settings for the GlobalProtect app.

1. From the Chromebook Management Console, select **Device management** > **Chrome management** > **App management**.

   The console displays the list of apps configured in all organization (org) units in your domain and displays the status of each app. Click an app **Status** to display the org units to which that status is applied.
2. Select the GlobalProtect app and then select **User settings**.

   If the app is not present, **SEARCH** for GlobalProtect in the Chrome Web Store.

© 2017 Palo Alto Networks, Inc.



STEP 2 | Configure policies and settings for everyone in an org unit.

1. Select the org unit where you want to configure settings and configure any of the following options:

   Selecting the top-level org unit applies settings to everyone in that unit; selecting a child org unit applies settings only to users within that child org unit.

   - **Allow installation**—Allow users install this app from the Chrome Web Store. By default, an org unit inherits the settings of its parent organization. To override the default settings, select **Inherit**, which toggles the **Override** setting.
   - **Force installation**—Install this app automatically and prevents users from removing it.
   - **Pin to taskbar**—If the app is installed, pin the app to the taskbar (in Chrome OS only).
   - **Add to Chrome Web Store collection**—Recommend this app to your users in the Chrome Web Store.
2. If you have not already done so, create a text file in JSON format that uses the following syntax and includes the FQDN or IP address of your GlobalProtect portal:

```
{
  "PortalAddress": {
    "Value": "192.0.2.191"
  }
}
```

3. On the **User settings** page, select **UPLOAD CONFIGURATION FILE** and then **Browse** to the GlobalProtect settings file.
4. **SAVE** your changes. Settings typically take effect within minutes, but it might take up to an hour to propagate through your organization.

STEP 3 | Test the connection.

After Chrome Management Console successfully deploys the app, Test the GlobalProtect app for Chrome OS

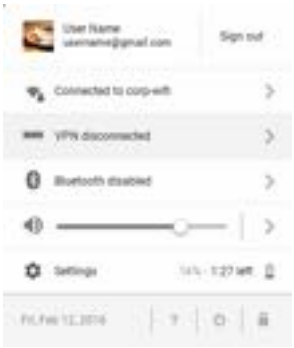
Test the GlobalProtect app for Chrome OS

Use the GlobalProtect app to view status and other information about the app or to collect logs, or reset the VPN connection settings. After you install and configure the app, it is not necessary to open the app to

© 2017 Palo Alto Networks, Inc.

establish a VPN connection. Instead, you can connect by selecting the portal from the VPN settings on the Chromebook.

STEP 1 | Log in to GlobalProtect.



1. Click the status area at the bottom right corner of the Chromebook.
2. Select **VPN disconnected** and then select the portal that you entered when configuring the GlobalProtect VPN settings.

   To view VPN settings before connecting, select the portal from **Settings** > **Private network**, and then click **Connect**.
3. Enter the **Username** and **Password** for the portal and click **Connect**. Repeat this step to enter the **Username** and **Password** for the gateway. If authentication is successful, GlobalProtect connects you to your corporate network. If enabled, the GlobalProtect welcome page will display.

STEP 2 | View the connection status. When the app is connected, the status area displays the VPN icon along the bottom of the Wi-Fi icon ( ).

- To view the portal to which you are connected, click the status area.



- To view additional information about the connection including the gateway to which you are connected, launch the GlobalProtect app. The main page displays connection information and (if applicable) any errors or warnings.

© 2017 Palo Alto Networks, Inc.

Deploy Agent Settings Transparently

As an alternative to deploying agent settings from the portal configuration, you can define them directly from the Windows registry or global Mac plist or—on Windows clients only—using the Windows Installer (Msiexec). The benefit is that it enables deployment of GlobalProtect agent settings to endpoints prior to their first connection to the GlobalProtect portal.

Settings defined in the portal configuration always override settings defined in the Windows registry or Mac plist. So if you define settings in the registry or plist, but the portal configuration specifies different settings, the settings the agent receives from the portal will override the settings defined on the client. This override also applies to login-related settings, such as whether to connect on-demand, whether to use single sign-on (SSO), and whether the agent can connect if the portal certificate is invalid. Therefore, you should avoid conflicting settings. In addition, the portal configuration is cached on the endpoint and that cached configuration is be used anytime the GlobalProtect agent is restarted or the client machine is rebooted.

The following sections describe the customizable agent settings available and how to deploy these settings transparently to Windows and Mac clients:

- Customizable Agent Settings on page 115
- Deploy Agent Settings to Windows Clients on page 122
- Deploy Agent Settings to Mac Clients on page 131



In addition to using Windows registry and Mac plist to deploy GlobalProtect agent settings, you can enable the GlobalProtect agent to collect specific Windows registry or Mac plist information from clients, including data on applications installed on the clients, processes running on the clients, and attributes or properties of those applications and processes. You can then monitor the data and add it to a security rule as matching criteria. Device traffic that matches registry settings you have defined can be enforced according to the security rule. Additionally, you can set up custom checks to Collect Application and Process Data From Clients *on page 184.*

Customizable Agent Settings

In addition to pre-deploying the portal address, you can also define the agent configuration settings. To Deploy Agent Settings to Windows Clients you define keys in the Windows registry (`HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect`), or, to Deploy Agent Settings to Mac Clients you define entries in the `PanSetup` dictionary of the Mac plist (`/Library/Preferences/com.paloaltonetworks.GlobalProtect.settings.plist`). On Windows clients only, you can also use the Windows Installer to Deploy Agent Settings from Msiexec.

The following topics describe each customizable agent setting. Settings defined in the GlobalProtect portal agent configuration take precedence over settings defined in the Windows registry or the Mac plist.



Some settings do not have a corresponding portal configuration settings on the web interface, and must be configured using Windows registry or Msiexec. These additional settings include: `can-prompt-user-credential`*,* `wrap-cp-guid`*, and* `filter-non-gpcp`*.*

- Agent Display Options
- User Behavior Options
- Agent Behavior Options
- Script Deployment Options

© 2017 Palo Alto Networks, Inc.

Agent Display Options

The following table lists the options that you can configure in the Windows registry and Mac plist to customize the display of the GlobalProtect agent.

Table 1: Table: Customizable Agent Settings

Portal Agent Configuration	Windows Registry/ Mac Plist	Msiexec Parameter	Default
Enable Advanced View	`enable-advanced-view yes | no`	`ENABLEADVANCEDVIEW="yes | no"`	`yes`
Display GlobalProtect Icon	`show-agent-icon yes | no`	`SHOWAGENTICON="yes | no"`	`yes`
Enable Rediscover Network Option	`rediscover-network yes | no`	`REDISCOVERNETWORK="yes | no"`	`yes`
Enable Resubmit Host Profile Option	`resubmit-host-info yes | no`	`RESUBMITHOSTINFO="yes | no"`	`yes`
Show System Tray Notifications	`show-system-tray-notifications yes | no`	`SHOWSYSTEMTRAYNOTIFICATIONS="yes | no"`	`yes`

User Behavior Options

The following table lists the options that you can configure in the Windows registry and Mac plist to customize how the user can interact with the GlobalProtect agent.

Table 2: Table: Customizable User Behavior Options

Portal Agent Configuration	Windows Registry/ Mac Plist	Msiexec Parameter	Default
Allow User to Change Portal Address	`can-change-portal yes | no`	`CANCHANGEPORTAL="yes | no"`	`yes`
Allow User to Dismiss Welcome Page	`enable-hide-welcome-page yes | no`	`ENABLEHIDEWELCOMEPAGE= "yes | no"`	`yes`
Allow User to Continue with Invalid Portal Server Certificate	`can-continue-if-portal-cert-invalid yes | no`	`CANCONTINUEIFPORTALCERTINVALID= "yes | no"`	`yes`
Allow User to Disable GlobalProtect App	`disable-allowed yes | no`	`DISABLEALLOWED="yes | no"`	`no`

© 2017 Palo Alto Networks, Inc.

Portal Agent Configuration	Windows Registry/ Mac Plist	Msiexec Parameter	Default
Save User Credentials Specify a 0 to prevent GlobalProtect from saving credentials, a 1 to save both username and password, or a 2 to save the username only.	`save-user-credentials 0 | 1 | 2`	`SAVEUSERCREDENTIALS 0 | 1 | 2`	n/a
Not in portal The **Allow user to save password** setting is deprecated in the web interface in PAN-OS 7.1 and later releases but is configurable from the Windows registry and Mac plist. Any value specified in the **Save User Credentials** field overwrites a value specified here.	`can-save-password yes | no`	`CANSAVEPASSWORD="yes | no"`	`yes`

Agent Behavior Options

The following table lists the options that you can configure in the Windows registry and Mac plist to customize the behavior of the GlobalProtect agent.

Table 3: Table: Customizable Agent Behavior Options

Portal Agent Configuration	Windows Registry/ Mac Plist	Msiexec Parameter	Default
Connect Method	`connect-method on-demand | pre-logon | user-logon`	`CONNECTMETHOD="on-demand | pre-logon | user-logon"`	`user-logon`
GlobalProtect App Config	`refresh-config-interval <hours>`	`REFRESHCONFIGINTERVAL= "<hours>"`	`24`

© 2017 Palo Alto Networks, Inc.

Portal Agent Configuration	Windows Registry/ Mac Plist	Msiexec Parameter	Default
Refresh Interval (hours)			
Update DNS Settings at Connect (Windows Only)	`flushdns yes | no`	`FLUSHDNS="yes | no"`	no
Send HIP Report Immediately if Windows Security Center (WSC) State Changes (Windows Only)	`wscautodetect yes | no`	`WSCAUTODETECT="yes | no"`	no
Detect Proxy for Each Connection (Windows Only)	`ProxyMultipleAuto Detection yes | no`	`ProxyMultipleAuto Detection="yes | no"`	no
Clear Single Sign-On Credentials on Logout (Windows Only)	`LogoutRemoveSSO yes | no`	`LogoutRemoveSSO="yes | no"`	yes
Use Default Authentication on Kerberos Authentication Failure (Windows Only)	`krb-auth-fail-fallback yes | no`	`KRBAUTHFAILFALLBACK= "yes | no"`	no
Custom Password Expiration Message (LDAP Authentication Only)	`PasswordExpiryMessage <message>`	`PasswordExpiryMessage "<message>"`	
Portal Connection Timeout (sec)	`PortalTimeout <portaltimeout>`	`PORTALTIMEOUT= "<portaltimeout>"`	30
TCP Connection Timeout (sec)	`ConnectTimeout <connecttimeout>`	`CONNECTTIMEOUT= "<connecttimeout>"`	60
TCP Receive Timeout (sec)	`ReceiveTimeout <receivetimeout>`	`RECEIVETIMEOUT= "<receivetimeout>"`	30
Client Certificate Store Lookup	`certificate-store-lookup user | machine | user and machine | invalid`	`CERTIFICATESTORELOOKUP= "user | machine | user and machine | invalid"`	user and machine

© 2017 Palo Alto Networks, Inc.

Portal Agent Configuration	Windows Registry/ Mac Plist	Msiexec Parameter	Default
SCEP Certificate Renewal Period (days)	**scep-certificate-renewal-period <renewalPeriod>**	n/a	**7**
Maximum Internal Gateway Connection Attempts	**max-internal-gateway-connection-attempts <maxValue>**	**MIGCA="<maxValue>"**	**0**
Extended Key Usage OID for Client Certificate	**ext-key-usage-oid-for-client-cert <oidValue>**	**EXTCERTOID="<oidValue>"**	n/a
User Switch Tunnel Rename Timeout (sec)	**user-switch-tunnel-rename-timeout <renameTimeout>**	n/a	**0**
Use Single Sign-On **(Windows Only)**	**use-sso yes | no**	**USESSO="yes | no"**	**yes**
Not in portal This setting specifies the default portal IP address (or hostname).	**portal <IPaddress>**	**PORTAL="<IPaddress>"**	n/a
Not in portal This setting enables GlobalProtect to initiate a VPN tunnel before a user logs in to the device and connects to the GlobalProtect portal.	**prelogon 1**	**PRELOGON="1"**	**1**
Windows only/ Not in portal This setting is used in conjunction with single sign-on (SSO) and indicates whether or not to prompt	**can-prompt-user-credential yes | no**	**CANPROMPTUSERCREDENTIAL="yes | no"**	**yes**

© 2017 Palo Alto Networks, Inc.

Portal Agent Configuration	Windows Registry/ Mac Plist	Msiexec Parameter	Default
the user for credentials if SSO fails.			
Windows only/ Not in portal This setting filters the third-party credential provider's tile from the Windows login page so that only the native Windows tile is displayed.*	**wrap-cp-guid {third party credential provider guid}**	**WRAPCPGUID="{guid_value]" FILTERNONGPCP="yes | no"**	**no**
Windows only/ Not in portal This setting is an additional option for the setting wrap-cp-guid, and allows the third-party credential provider tile to be displayed on the Windows login page, in addition to the native Windows logon tile.*	**filter-non-gpcp no**	n/a	n/a

*For detailed steps to enable these settings using the Windows registry or Windows Installer (Msiexec), see SSO Wrapping for Third-Party Credential Providers on Windows Clients.

Script Deployment Options

The following table displays options that enable GlobalProtect to initiate scripts before and after establishing a VPN tunnel and before disconnecting a VPN tunnel. Because these options are not available in the portal, you must define the values for the relevant key—either pre-vpn-connect, post-vpn-connect, or pre-vpn-disconnect—from the Windows registry or Mac plist. For detailed steps to deploy scripts, see Deploy Scripts Using the Windows Registry, Deploy Scripts Using Msiexec, or Deploy Scripts Using the Mac Plist.

Table: Customizable Script Deployment Options

© 2017 Palo Alto Networks, Inc.

Portal Agent Configuration	Windows Registry/ Mac Plist	Msiexec Parameter	Default
Execute the script specified in the command setting (including any parameters passed to the script).  *Environmental variables are supported.*  *Specify the full path in commands.*	**command <parameter1> <parameter2> [...]** Windows example: **command %userprofile%\vpn_script.bat c:test_user** Mac example: **command $HOME/vpn_script.sh /Users/test_user test_user**	**PREVPNCONNECTCOMMAND="<parameter1> <parameter2> [...]"** **POSTVPNCONNECTCOMMAND="<parameter1> <parameter2> [...]"** **PREVPNDISCONNECTCOMMAND="<parameter1> <parameter2> [...]"**	**n/a**
(Optional) Specify the privileges under which the command(s) can run (default is user: if you do not specify the context, the command runs as the current active user).	**context admin | user**	**PREVPNCONNECTCONTEXT="admin | user"** **POSTVPNCONNECTCONTEXT="admin | user"** **PREVPNDISCONNECTCONTEXT="admin | user"**	**user**
(Optional) Specify the number of seconds the GlobalProtect client waits for the command to execute (range is 0-120). If the command does not complete before the timeout, the client proceeds to establish or disconnect from the VPN tunnel. A value of 0 (the default) means the client will not wait to execute the command.  *Not supported for post-vpn-connect.*	**timeout <value>** Example: **timeout 60**	**PREVPNCONNECTTIMEOUT="<value>"** **POSTVPNCONNECTTIMEOUT="<value>"** **PREVPNDISCONNECTTIMEOUT="<value>"**	**0**
(Optional) Specify the full path of a file used in a command. The GlobalProtect client will verify the integrity of the file by checking it against	**file <path_file>**	**PREVPNCONNECTFILE="<path_file>"** **POSTVPNCONNECTFILE="<path_file>"** **PREVPNDISCONNECTFILE="<path_file>"**	**n/a**

© 2017 Palo Alto Networks, Inc.

Portal Agent Configuration	Windows Registry/ Mac Plist	Msiexec Parameter	Default
the value specified in the checksum key.  *Environmental variables are supported.*			
(Optional) Specify the sha256 checksum of the file referred to in the file key. If the checksum is specified, the GlobalProtect client executes the command(s) only if the checksum generated by the GlobalProtect client matches the checksum value specified here.	**checksum <value>**	**PREVPNCONNECTCHECKSUM= "<value>"** **POSTVPNCONNECTCHECKSUM= "<value>"** **PREVPNDISCONNECTCHECKSUM ="<value>"**	**n/a**
(Optional) Specify an error message to inform the user that the command(s) cannot execute or if the command(s) exited with a non-zero return code.  *The message must be 1,024 or fewer ANSI characters.*	**error-msg <message>** Example: **error-msg Failed executing pre-vpn-connect action!**	**PREVPNCONNECTERRORMSG= "<message>"** **POSTVPNCONNECTERRORMSG= "<message>"** **PREVPNDISCONNECTERRORMSG ="<message>"**	**n/a**

Deploy Agent Settings to Windows Clients

Use Windows registry or the Windows Installer (Msiexec) to deploy the GlobalProtect agent and settings to Windows clients transparently.

- Deploy Agent Settings in the Windows Registry
- Deploy Agent Settings from Msiexec
- Deploy Scripts Using the Windows Registry
- Windows OS Batch Script Examples
- Deploy Scripts Using Msiexec
- SSO Wrapping for Third-Party Credential Providers on Windows Clients
- Enable SSO Wrapping for Third-Party Credentials with the Windows Registry
- Enable SSO Wrapping for Third-Party Credentials with the Windows Installer

© 2017 Palo Alto Networks, Inc.

Deploy Agent Settings in the Windows Registry

You can enable deployment of GlobalProtect agent settings to Windows clients prior to their first connection to the GlobalProtect portal by using the Windows registry. Use the options described in the following table to begin using the Windows registry to customize agent settings for Windows clients.

In addition to using Windows registry to deploy GlobalProtect agent settings, you can enable the GlobalProtect agent to collect specific Windows registry information from Windows clients. You can then monitor the data and add it to a security rule as matching criteria. Device traffic that matches registry settings you have defined can be enforced according to the security rule. Additionally, you can set up custom checks to Collect Application and Process Data From Clients.

- Locate the GlobalProtect agent customization settings in the Windows registry.

  Open the Windows registry (enter **`regedit`** at the command prompt) and go to:

  `HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\Settings\`

- Set the portal name.

  If you do not want the user to manually enter the portal address even for the first connection, you can pre-deploy the portal address through the Windows registry: (`HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\PanSetup` with key `Portal`).



- Deploy various settings to the Windows client from the Windows registry, including configuring the connect method for the GlobalProtect agent and enabling single sign-on (SSO).

  View Customizable Agent Settings for a full list of the commands and values you can set up using the Windows registry.

- Enable the GlobalProtect agent to wrap third-party credentials on the Windows client, allowing for SSO when using a third-party credential provider.

  Enable SSO Wrapping for Third-Party Credentials with the Windows Registry.

© 2017 Palo Alto Networks, Inc.

Deploy Agent Settings from Msiexec

On Windows endpoints, you have the option to deploy the agent and the settings automatically from the Windows Installer (Msiexec) by using the following syntax:

```
msiexec.exe /i GlobalProtect.msi <SETTING>="<value>"
```

Msiexec is an executable program that installs or configures a product from the command line. On systems running Microsoft Windows XP or a later OS, the maximum length of the string that you can use at the command prompt is 8,191 characters.

Msiexec Example	Description
`msiexec.exe /i GlobalProtect.msi /quiet PORTAL="portal.acme.com"`	Install GlobalProtect in quiet mode (no user interaction) and configure the portal address.
`msiexec.exe /i GlobalProtect.msi CANCONTINUEIFPORTALCERTINVALID="no"`	Install GlobalProtect with the option to prevent users from connecting to the portal if the certificate is not valid.

For a complete list of settings and the corresponding default values, see Customizable Agent Settings.

To set up the GlobalProtect agent to wrap third-party credentials on a Windows client from Msiexec, see Enable SSO Wrapping for Third-Party Credentials with the Windows Installer.

Deploy Scripts Using the Windows Registry

You can enable deployment of custom scripts to Windows endpoints using the Windows registry.

You can configure the GlobalProtect agent to initiate and run a script for any or all of the following events: before and after establishing the tunnel, and before disconnecting the tunnel. To run the script at a particular event, reference the batch script from a command registry entry for that event.

Depending on the configuration settings, the GlobalProtect agent can run a script before and after the agent establishes a VPN tunnel with the gateway, and before the agent disconnects from the VPN tunnel. Use the following workflow to get started using the Windows registry to customize agent settings for Windows clients.

The registry settings that enable you to deploy scripts are supported in GlobalProtect clients running GlobalProtect agent 2.3 and later releases.

STEP 1 | Open the Windows registry, and locate the GlobalProtect agent customization settings.

Open the Windows registry (enter **regedit** in the command prompt) and go to the location of the key depending on when you want to execute scripts (pre/post connect or pre disconnect):

```
HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\Settings\pre-vpn-connect
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\Settings\post-vpn-connect
```


© 2017 Palo Alto Networks, Inc.

`HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\Settings\pre-vpn-disconnect`

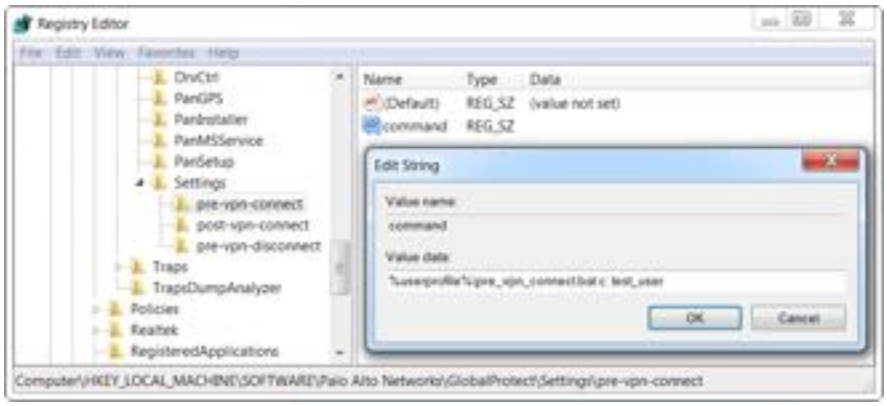
If the key does not exist within the Settings key, create it (right-click Settings and select New > Key).

STEP 2 | Enable the GlobalProtect agent to run scripts by creating a new String Value named **command**.

The batch file specified here should contain the specific script (including any parameters passed to the script) that you want run on the device. For examples, see Windows OS Batch Script Examples.

1. If the `command` string does not already exist, create it (right-click the `pre-vpn-connect`, `post-vpn-connect`, or `pre-vpn-disconnect` key, select **New** > **String Value**, and name it **command**).
2. Right click `command` and select **Modify**.
3. Enter the commands or script that the GlobalProtect agent should run. For example:

```
%userprofile%\pre_vpn_connect.bat c:
test_user
```



STEP 3 | (**Optional**) Add additional registry entries as needed for each command.

Create or modify registry strings and their corresponding values, including `context`, `timeout`, `file`, `checksum`, or `error-msg`. For additional information, see Customizable Agent Settings.

Windows OS Batch Script Examples

You can configure the GlobalProtect agent to initiate and run a script for any or all of the following events: before and after establishing the tunnel, and before disconnecting the tunnel. To run the script at a particular event, reference the batch script from a command registry entry for that event. The following examples show scripts you can run on Windows systems at pre-connect, post-connect, and pre-disconnect events:

Example: Exclude Traffic from the VPN Tunnel on Windows Endpoints

To exclude traffic from the VPN tunnel after establishing the VPN connection, reference the following script from a `command` registry entry for a post-vpn-connect event. This enables you to selectively exclude routes and to send all other traffic through the VPN tunnel.

As a best practice, delete any exclude network routes that were previously added before adding the new exclude routes. In most cases, when a user moves between networks (such as when switching between Wi-Fi and a local network) the old network routes are

© 2017 Palo Alto Networks, Inc.

automatically deleted. In the event that the old network routes persist, following this best practice ensures that traffic destined for the exclude routes will go through the gateway of the new network instead of the gateway of the old network.



For a script that you can copy and paste, go here.

```
@echo off
REM Run this script (route_exclude) post-vpn-connect.
REM Add exclude routes. This allows traffic to these
network and hosts to go directly and not use the tunnel.
REM Syntax: route_exclude <network1> <mask1> <network2>
<mask2> ...<networkN> <maskN>
REM Example-1: route_exclude 10.0.0.0 255.0.0.0
REM Example-2: route_exclude 10.0.0.0 255.0.0.0 192.168.17.0
255.255.255.0
REM Example-3: route_exclude 10.0.0.0 255.0.0.0 192.168.17.0
255.255.255.0 192.168.24.25 255.255.255.255

REM Initialize 'DefaultGateway'
set "DefaultGateway="

REM Use the route print command and find the DefaultGateway
on the endpoint
@For /f "tokens=3" %%* in (
   'route.exe print ^|findstr "\<0.0.0.0\>"'
   ) Do if not defined DefaultGateway Set "DefaultGateway=%%*"

REM Use the route add command to add the exclude routes
:add_route
if "%1" =="" goto end
route delete %1
route add %1 mask %2 %DefaultGateway%
shift
shift
goto add_route
:end
```

Example: Mount a Network Share on Windows Endpoints

To mount a network share after establishing a VPN connection, reference the following script from a `command` registry entry for a post-vpn-connect event:

```
@echo off
REM Mount filer1 to Z: drive
net use Z: \\filer1.mycompany.local\share /user:mycompany\user1
```

Deploy Scripts Using Msiexec

On Windows clients, you can use the Windows Installer (Msiexec) to deploy the agent, agent settings, and scripts that the agent will run automatically (see Customizable Agent Settings). To do so, use the following syntax:

```
msiexec.exe /i GlobalProtect.msi <SETTING>="<value>"
```


© 2017 Palo Alto Networks, Inc.

Msiexec is an executable program that installs or configures a product from a command line. On systems running Microsoft Windows XP or a later release, the maximum length of the string that you can use at the command prompt is 8,191 characters.

This limitation applies to the command line, individual environment variables (such as the USERPROFILE variable) that are inherited by other processes, and all environment variable expansions. If you run batch files from the command line, this limitation also applies to batch file processing.

For example, to deploy scripts that run at specific connect or disconnect events, you can use syntax similar to the following examples:

Example: Use Msiexec to Deploy Scripts that Run Before a Connect Event

For a script that you can copy and paste, go here.

```
msiexec.exe /i GlobalProtect.msi
PREVPNCONNECTCOMMAND="%userprofile%\pre_vpn_connect.bat c: test_user"
PREVPNCONNECTCONTEXT="user"
PREVPNCONNECTTIMEOUT="60"
PREVPNCONNECTFILE="C:\Users\test_user\pre_vpn_connect.bat"
PREVPNCONNECTCHECKSUM="a48ad33695a44de887bba8f2f3174fd8fb01a46a19e3ec9078b011
8647ccf599"
PREVPNCONNECTERRORMSG="Failed executing pre-vpn-connect action."
```

For a complete list of settings and the corresponding default values, see Customizable Agent Settings. Or, for examples of batch scripts, see Windows OS Batch Script Examples.

Example: Use Msiexec to Deploy Scripts that Run at Pre-Connect, Post-Connect, and Pre-Disconnect Events

For a script that you can copy and paste, go here.

```
msiexec.exe /i GlobalProtect.msi
PREVPNCONNECTCOMMAND="%userprofile%\pre_vpn_connect.bat c: test_user"
PREVPNCONNECTCONTEXT="user"
PREVPNCONNECTTIMEOUT="60"
PREVPNCONNECTFILE="C:\Users\test_user\pre_vpn_connect.bat"
PREVPNCONNECTCHECKSUM="a48ad33695a44de887bba8f2f3174fd8fb01a46a19e3ec9078b011
8647ccf599"
PREVPNCONNECTERRORMSG="Failed executing pre-vpn-connect action."
POSTVPNCONNECTCOMMAND="c:\users\test_user\post_vpn_connect.bat c: test_user"
POSTVPNCONNECTCONTEXT="admin"
POSTVPNCONNECTFILE="%userprofile%\post_vpn_connect.bat"
POSTVPNCONNECTCHECKSUM="b48ad33695a44de887bba8f2f3174fd8fb01a46a19e3ec9078b011
8647ccf598"
POSTVPNCONNECTERRORMSG="Failed executing post-vpn-connect action."
PREVPNDISCONNECTCOMMAND="%userprofile%\pre_vpn_disconnect.bat c: test_user"
PREVPNDISCONNECTCONTEXT="admin"
PREVPNDISCONNECTTIMEOUT="0"
PREVPNDISCONNECTFILE="C:\Users\test_user\pre_vpn_disconnect.bat"
PREVPNDISCONNECTCHECKSUM="c48ad33695a44de887bba8f2f3174fd8fb01a46a19e3ec9078b0
118647ccf597"
```


© 2017 Palo Alto Networks, Inc.

```
PREVPNDISCONNECTERRORMSG="Failed executing pre-vpn-disconnect action."
```

For a complete list of settings and the corresponding default values, see Customizable Agent Settings. Or, for examples of batch scripts, see Windows OS Batch Script Examples.

SSO Wrapping for Third-Party Credential Providers on Windows Clients

On Windows 7 and Windows Vista clients, the GlobalProtect agent utilizes the Microsoft credential provider framework to support single sign-on (SSO). With SSO, the GlobalProtect credential provider wraps the Windows native credential provider, which enables GlobalProtect to use Windows login credentials to automatically authenticate and connect to the GlobalProtect portal and gateway.

In some scenarios when other third-party credential providers also exist on the client, the GlobalProtect credential provider is unable to gather a user's Windows login credentials and, as a result, GlobalProtect fails to automatically connect to the GlobalProtect portal and gateway. If SSO fails, you can identify the third-party credential provider and then configure the GlobalProtect agent to wrap those third-party credentials, which enables users to successfully authenticate to Windows, GlobalProtect, and the third-party credential provider—all in a single step—using only their Windows login credentials when they log in to their Windows system.

Optionally, you can configure Windows to display separate login tiles: one for each third-party credential provider and another for the native Windows login. This is useful when a third-party credential provider adds additional functionality in the login tile that does not apply to GlobalProtect.

Use the Windows registry or the Windows Installer (Msiexec) to allow GlobalProtect to wrap third-party credentials:

- Enable SSO Wrapping for Third-Party Credentials with the Windows Registry on page 128
- Enable SSO Wrapping for Third-Party Credentials with the Windows Installer on page 130



GlobalProtect SSO wrapping for third-party credential providers (CPs) is dependent on the third-party CP settings and, in some cases, GlobalProtect SSO wrapping might not work correctly if the third-party CP implementation does not allow GlobalProtect to successfully wrap their CP.

Enable SSO Wrapping for Third-Party Credentials with the Windows Registry

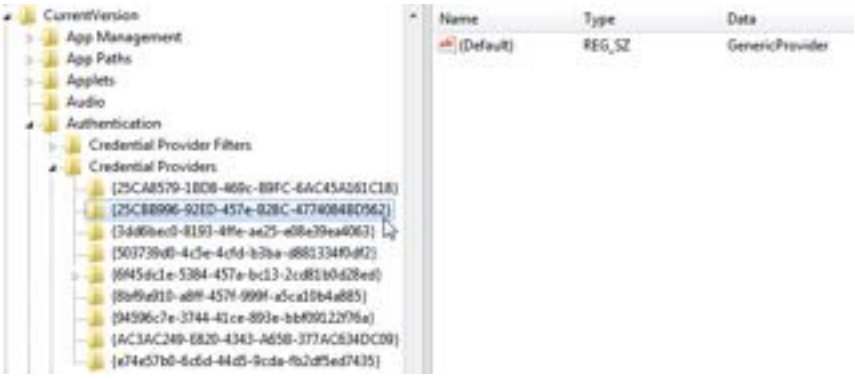
Use the following steps in the Windows registry to enable SSO to wrap third-party credentials on Windows 7 and Windows Vista clients.

STEP 1 | Open the Windows registry and locate the globally unique identifier (GUID) for the third-party credential provider that you want to wrap.

1. From the command prompt, enter the command **`regedit`** to open the Windows registry.
2. Locate currently installed credential providers at the following location:

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\ CurrentVersion
   \Authentication\Credential Providers.
   ```
3. Copy the GUID key for the credential provider that you want to wrap (including the curly brackets—`{` and `}` —on either end of the GUID):

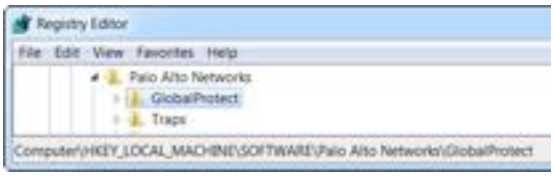
© 2017 Palo Alto Networks, Inc.



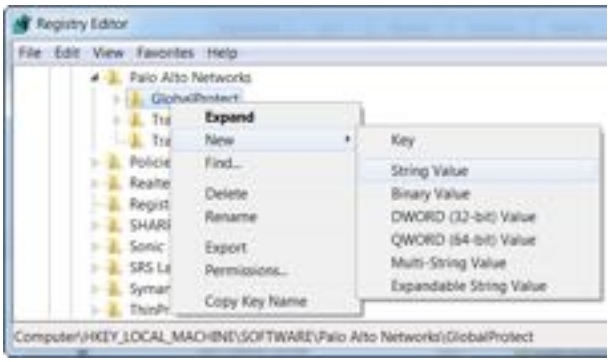
STEP 2 | Enable SSO wrapping for third-party credential providers by adding the setting **`wrap-cp-guid`** to the GlobalProtect registry.

1. Go to the following Windows registry location:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\ GlobalProtect:
```



2. Add a new **String Value**:



3. Enter values for the **String Value**:
   - **Name**: **`wrap-cp-guid`**
   - **Value data**: {<**`third-party credential provider GUID`**>}

   For the Value data field, the GUID value that you enter must be enclosed with curly brackets: `{` and `}`.

   The following is an example of what a third-party credential provider GUID in the **Value data** field might look like:

```
{A1DA9BCC-9720-4921-8373-A8EC5D48450F}
```

For the new String Value, `wrap-cp-guid` is displayed as the String Value's Name and the GUID is displayed as the Data.

© 2017 Palo Alto Networks, Inc.



STEP 3 | Next Steps:

- You can configure SSO wrapping for third-party credential providers successfully by completing steps 1 and 2. With this setup, the native Windows logon tile is displayed to users. Users click the tile and log in to the system with their Windows credentials and that single login authenticates the users to Windows, GlobalProtect, and the third-party credential provider.
- (**Optional**) If you want to display two tiles to users at login, the native Windows tile and the tile for the third-party credential provider, continue to 4.

STEP 4 | (**Optional**) Allow the third-party credential provider tile to be displayed to users at login.

Add a second **String Value** with the **Name** **`filter-non-gpcp`** and enter **`no`** for the string's **Value data**:



With this string value added to the GlobalProtect settings, two login options are presented to users when logging in to their Windows system: the native Windows tile and the third-party credential provider's tile.

Enable SSO Wrapping for Third-Party Credentials with the Windows Installer

Use the following options in the Windows Installer (Msiexec) to enable SSO to wrap third-party credential providers on Windows 7 and Windows Vista clients.

- Wrap third-party credentials and display the native tile to users at login. Users click the tile and log in to the system with their native Windows credentials and that single login authenticates users to Windows, GlobalProtect, and the third-party credential provider.

  Use the following syntax from the Windows Installer (Msiexec):

```
msiexec.exe /i GlobalProtect.msi WRAPCPGUID="{guid_value}"
 FILTERNONGPCP="yes"
```

  In the syntax above, the **`FILTERNONGPCP`** parameter simplifies authentication for the user by filtering the option to log in to the system using the third-party credentials.

- If you would like users to have the option to log in with the third-party credentials, use the following syntax from the Msiexec:

```
msiexec.exe /i GlobalProtect.msi WRAPCPGUID="{guid_value}"
FILTERNONGPCP="no"
```

  In the syntax above, the **`FILTERNONGPCP`** parameter is set to **`"no"`**, which filters out the third-party credential provider's logon tile so that only the native tile displays. In this case, both the native Windows tile and the third-party credential provider tile is displayed to users when logging in to the Windows system.

© 2017 Palo Alto Networks, Inc.

Deploy Agent Settings to Mac Clients

Use the Mac global plist (property list) file to set GlobalProtect agent customization settings for or to deploy scripts to Mac endpoints.

- Deploy Agent Settings in the Mac Plist
- Deploy Scripts Using the Mac Plist
- Mac OS Script Examples

Deploy Agent Settings in the Mac Plist

You can set the GlobalProtect agent customization settings in the Mac global plist (Property list) file. This enables deployment of GlobalProtect agent settings to Mac endpoints prior to their first connection to the GlobalProtect portal.

On Mac systems, plist files are either located in `/Library/Preferences` or in `~/Library/Preferences`. The tilde (~) symbol indicates that the location is in the current user's home folder. The GlobalProtect agent on a Mac client first checks for the GlobalProtect plist settings. If the plist does not exist at that location, the GlobalProtect agent searches for plist settings in `~/Library/Preferences`.



In addition to using the Mac plist to deploy GlobalProtect agent settings, you can enable the GlobalProtect agent to collect specific Mac plist information from clients. You can then monitor the data and add it to a security rule as matching criteria. Device traffic that matches registry settings you have defined can be enforced according to the security rule. Additionally, you can set up custom checks to Collect Application and Process Data From Clients.

- Open the GlobalProtect plist file and locate the GlobalProtect agent customization settings.

  Use Xcode or an alternate plist editor to open the plist file:

  `/Library/Preferences/com.paloaltonetworks.GlobalProtect.settings.plist`

  Then go to:

  `/Palo Alto Networks/GlobalProtect/Settings`

  If the `Settings` dictionary does not exist, create it. Then add each key to the Settings dictionary as a string.

- Set the portal name.

  If you don't want the user to manually enter the portal address even for the first connection, you can pre-deploy the portal address through the Mac plist. Under the `PanSetup` dictionary, configure an entry for `Portal`.

- Deploy various settings to the Mac client from the Mac plist, including configuring the connect method for the GlobalProtect agent.

  View Customizable Agent Settings for a full list of the keys and values that you can configure using the Mac plist.

Deploy Scripts Using the Mac Plist

When a user connects to the GlobalProtect gateway for the first time, the GlobalProtect agent downloads a configuration file and stores agent settings in a GlobalProtect Mac property file (plist). In addition to making

© 2017 Palo Alto Networks, Inc.

changes to the agent settings, you use the Mac plist to deploy scripts at any or all of the following events: before and after establishing the tunnel, and before disconnecting the tunnel. Use the following workflow to get started using the Mac plist to deploy scripts to Mac endpoints.

The Mac plist settings that enable you to deploy scripts are supported in GlobalProtect clients running GlobalProtect agent 2.3 and later releases.

STEP 1 | (**Clients running Mac OS X 10.9 or a later OS**) Flush the settings cache. This prevents the OS from using the cached preferences after making changes to the plist.

To clear the default preferences cache, run the **`killall cfprefsd`** command from a Mac terminal.

STEP 2 | Open the GlobalProtect plist file, and locate or create the GlobalProtect dictionary associated with the connect or disconnect event. The dictionary under which you will add the settings will determine when the GlobalProtect agent runs the script(s).

Use Xcode or an alternate plist editor to open the plist file (`/Library/Preferences/com.paloaltonetworks.GlobalProtect.settings.plist`) and go to the location of the dictionary:

- `/Palo Alto Networks/GlobalProtect/Settings/pre-vpn-connect`
- `/Palo Alto Networks/GlobalProtect/Settings/post-vpn-connect`
- `/Palo Alto Networks/GlobalProtect/Settings/pre-vpn-disconnect`

If `Settings` dictionary does not exist, create it. Then, in `Settings`, create a new dictionary for the event or events at which you want to run scripts.

STEP 3 | Enable the GlobalProtect agent to run scripts by creating a new **`String`** named **`command`**.

The value specified here should reference the shell script (and the parameters to pass to the script) that you want run on your devices. See Mac OS Script Examples on page 133.

If the `command` string does not already exist, add it to the dictionary and specify the script and parameters in the **Value** field, for example:

```
$HOME\pre_vpn_connect.sh
/Users/username username
```

Environmental variables are supported.

As a best practice, specify the full path in commands.

STEP 4 | (**Optional**) Add additional settings related to the command, including administrator privileges, a timeout value for the script, checksum value for the batch file, and an error message to display if the command fails to execute successfully.

Create or modify additional strings in the plist (`context`, `timeout`, `file`, `checksum`, and/or `error-msg`) and enter their corresponding values. For additional information, see Customizable Agent Settings on page 115.

STEP 5 | Save the changes to the plist file.

Save the plist.

© 2017 Palo Alto Networks, Inc.

Mac OS Script Examples

You can configure the GlobalProtect agent to initiate and run a script for any or all of the following events: before and after establishing the tunnel, and before disconnecting the tunnel. To run the script at a particular event, reference the shell script from a `command` plist entry for that event. The following topics show examples of scripts that you can run at pre-connect, post-connect and pre-disconnect events:

Example: Terminate All Established SSH Sessions on Mac Endpoints

To force termination of all established SSH sessions before setting up the VPN tunnel, reference the following script from a `command` plist entry for a pre-vpn-connect event. Similarly, you can re-establish the sessions after establishing the GlobalProtect VPN tunnel by using a script that you reference from the `command` plist entry for a post-vpn-connect event. This can be useful if you want to force all SSH traffic to traverse the GlobalProtect VPN tunnel.

```
#!bin/bash
# Identify all SSH sessions and force kill them
ps | grep ssh | grep -v grep | awk '{ print $1 }' | xargs
kill -9
```

Example: Mount a Network Share on Mac Endpoints

To mount a network share after establishing a VPN connection, reference the following script from a `command` plist entry for a post-vpn-connect event:



For a script that you can copy and paste, go here.

```
#!/bin/bash
mkdir $1
mount -t smbfs //username:password@10.101.2.17/shares/Departments/Engineering/
SW_eng/username/folder
$1
sleep 1
```


© 2017 Palo Alto Networks, Inc.

GlobalProtect Clientless VPN

GlobalProtect Clientless VPN provides secure remote access to common enterprise web applications. Users have the advantage of secure access from SSL-enabled web browsers without installing GlobalProtect client software. This is useful when you need to enable partner or contractor access to applications, and to safely enable unmanaged assets, including personal devices. You can configure the GlobalProtect portal landing page to provide access to web applications based on users and user groups and also allow single-sign on to SAML-enabled applications. The following topics provide information on how to configure and troubleshoot Clientless VPN.

> Clientless VPN Overview
> Supported Technologies
> Configure Clientless VPN
> Troubleshoot Clientless VPN

© 2017 Palo Alto Networks, Inc.

Clientless VPN Overview

When you configure Clientless VPN, remote users can log in to the GlobalProtect portal using a web browser and launch the web applications you publish for the user. Based on users or user groups, you can allow users to access a set of applications that you make available to them, or allow them to access additional corporate applications by entering a custom application URL.

After logging in to the portal, users see a published applications page with a list of web applications they can launch. (You can use the default applications landing page on the GlobalProtect portal or create a custom landing page for your enterprise.)

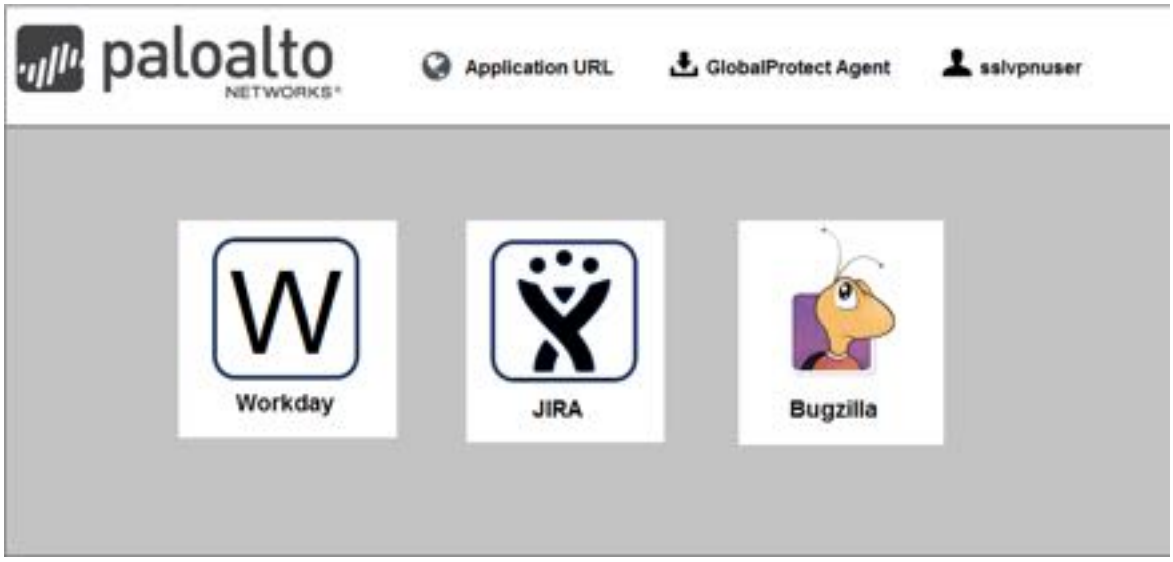


Figure 1: Applications Landing Page for Clientless VPN

Because this page replaces the default portal landing page, it includes a link to the GlobalProtect agent download page. If configured, users can also select **Application URL** and enter URLs to launch additional, unpublished corporate web applications.

When you configure only one web application (and disable access to unpublished applications), instead of taking the user to the published applications page, the application will launch automatically as soon as the user logs in. If you do not configure GlobalProtect Clientless VPN, users will see the agent software download page when they log in to the portal.

When you configure GlobalProtect Clientless VPN, you need security policies to allow traffic from GlobalProtect endpoints to the security zone associated with the GlobalProtect portal that hosts the published applications landing page and security policies to allow user-based traffic from the GlobalProtect portal zone to the security zone where the published application servers are hosted. The security policies you define control which users have permission to use each published application.

© 2017 Palo Alto Networks, Inc.

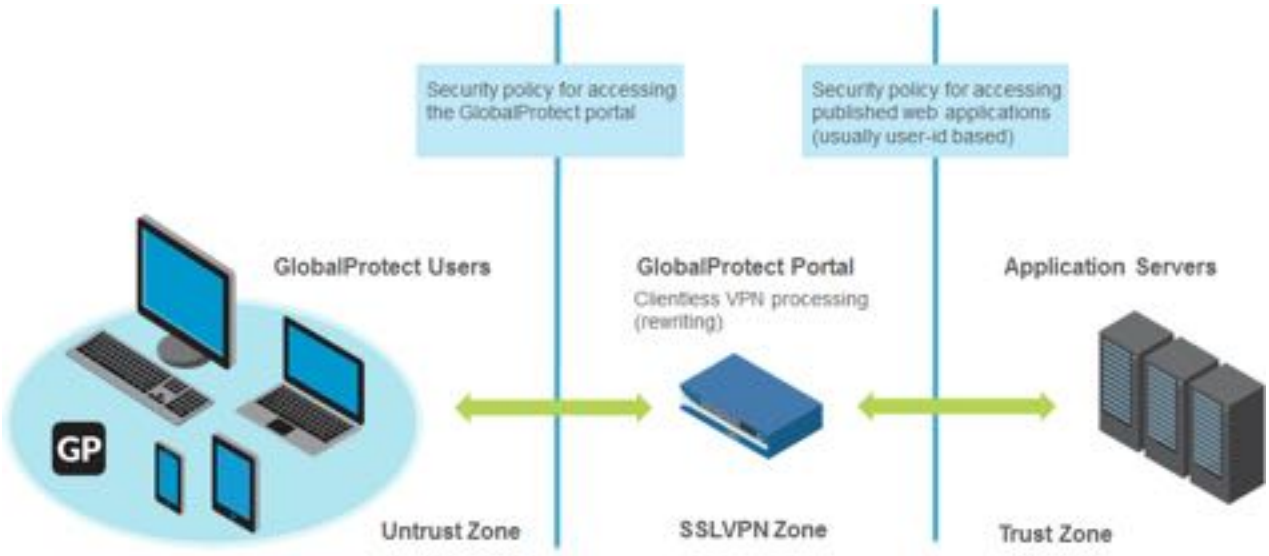


Figure 2: Zones and Security Policy for Clientless VPN

© 2017 Palo Alto Networks, Inc.

Supported Technologies

You can configure the GlobalProtect portal to provide secure remote access to common enterprise web applications. For best results, make sure you throughly test your Clientless VPN applications in a controlled environment before deploying them or making them available to a large number of users.

Technology	Supported Version

Web application technologies	• HTML
	• HTML5
	• HTML5-Web-Sockets
	• Javascript
	• Remote desktop protocol (RDP), VNC, or SSH through third-party HTML middleware applications—Examples of third-party middleware applications that use supported web technologies include:
	  • HOBLink WebTerm—Uses HTML5
	  • VMware Horizon, vSphere, vCenter—Uses HTML5
	  • Apache Guacamole—Uses HTML5
	• Adobe Flash—With Clientless VPN, browsers can display flash content or files like Microsoft Word documents or Adobe PDFs; however, Clientless VPN cannot rewrite HTML content links inside the content to route the links through the tunnel.
	Other technologies (such as Microsoft Silverlight or XML/XSLT) are not supported.

Operating systems	• Windows • Mac • iOS • Android • Chrome • Linux
Supported browsers	• Chrome • Internet Explorer • Safari • Firefox

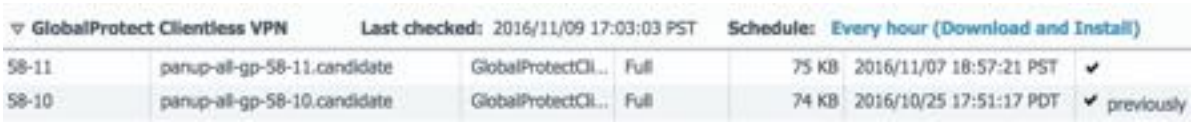
© 2017 Palo Alto Networks, Inc.

Configure Clientless VPN

To configure GlobalProtect Clientless VPN:

STEP 1 | Before you begin:

- Install a GlobalProtect subscription on the firewall that hosts the Clientless VPN from the GlobalProtect portal. Refer to Active Licenses and Subscriptions.
- Install the GlobalProtect Clientless VPN dynamic update (see Install Content and Software Updates).



- As a best practice, configure a separate FQDN for the GlobalProtect portal that hosts Clientless VPN. Do not use the same FQDN as the PAN-OS Web Interface.
- Host the GlobalProtect portal on the standard SSL port (TCP port 443). Non-standard ports are not supported.

STEP 2 | Configure the applications that are available using GlobalProtect Clientless VPN. The GlobalProtect portal displays these applications on the landing page that users see when they log in (the applications landing page).

1. Select **Network** > **GlobalProtect** > **Clientless Apps** and **Add** one or more applications. For each application, specify the following:
   - **Name**—A descriptive name for the application (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
   - **Location** (for a firewall that is in multiple virtual system mode)—the virtual system (vsys) where the Clientless VPN applications are available. For a firewall that is not in multi-vsys mode, the **Location** field does not appear.
   - **Application Home URL**—The URL where the web application is located (up to 4095 characters).
   - **Application Description** (**Optional**)—A brief description of the application (up to 255 characters).
   - **Application Icon** (**Optional**)—An icon to identify the application on the published application page. You can browse to upload the icon.
2. Click **OK**.

STEP 3 | (**Optional**). Create groups to manage sets of web applications.

Clientless App Groups are useful if you want to manage multiple collections of applications and provide access based on user groups. For example, financial applications for the G&A team or developer applications for the Engineering team.

1. Select **Network** > **GlobalProtect** > **Clientless App Groups**. **Add** a new Clientless VPN application group and specify the following:
   - **Name**—A descriptive name for the application group (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
   - **Location** (for a firewall that is in multiple virtual system mode)—the virtual system (vsys) where the Clientless VPN application group is available. For a firewall that is not in multi-vsys mode, the **Location** field does not appear.
2. In the **Applications** area, **Add** applications to the group. You can select from the list of existing Clientless VPN applications or define a **New Clientless App**.
3. Click **OK**.

© 2017 Palo Alto Networks, Inc.

STEP 4 | Configure the GlobalProtect portal to provide the Clientless VPN service.

1. Select **Network** > **GlobalProtect** > **Portal** and select an existing portal configuration or **Add** a new portal. Refer to Set Up Access to the GlobalProtect Portal.
2. In the **Authentication** tab, you can:
   - (**Optional**) Create a new client authentication specifically for Clientless VPN. In this case, choose **Browser** as the **OS** for **Client Authentication**.
   - Use an existing client authentication.
3. In **Clientless** > **General**, select **Clientless VPN** to enable the portal service and configure the following:
   - Specify a **Hostname** (IP address or fully-qualified domain name) for the GlobalProtect portal that hosts the applications landing page. This hostname is used for rewriting application URLs. (For more information on URL rewriting, refer to 8).

     If you use Network Address Translation (NAT) to provide access to the GlobalProtect portal, the IP address or FQDN you enter must match (or resolve to) the NAT IP address for the GlobalProtect portal (the public IP address). Because users cannot access the GlobalProtect portal on a custom port, the pre-NAT port must also be TCP port 443.

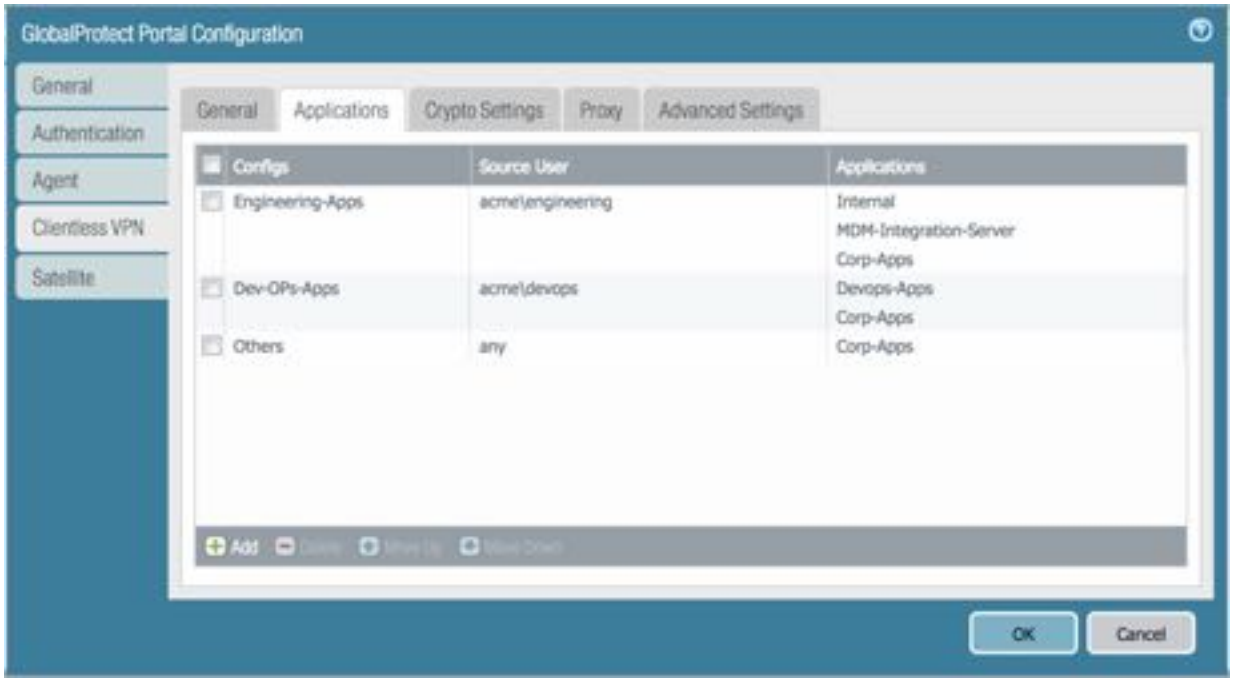
   - Specify a **Security Zone**. This zone is used as a source zone for the traffic between the firewall and the applications. Security rules defined from this zone to the application zone determine which applications can be accessed.
   - Select a **DNS Proxy** server or configure a **New DNS Proxy**. GlobalProtect will use this proxy to resolve application names. Refer to DNS Proxy Object.
   - **Login Lifetime**—Specify the maximum hours or minutes that a Clientless VPN session is valid. The typical session time is 3 hours. The range for hours is 1-24; the range for minutes is 60-1440. After the session expires, users must re-authenticate and start a new Clientless VPN session.
   - **Inactivity Timeout**—Specify the number of hours or minutes that a Clientless VPN session can remain idle. The typical inactivity timeout is 30 minutes. The range for hours is 1-24; the range for minutes is 5 to 1440. If there is no user activity during the specified amount of time, users must re-authenticate and start a new Clientless VPN session.
   - **Max User**—Specify the maximum number of users who can be logged into the portal at the same time. If no value is specified, then endpoint capacity is assumed. If the endpoint capacity is unknown, then a capacity of 50 users is assumed. When the maximum number of users is reached, additional Clientless VPN users cannot log in to the portal.

STEP 5 | Map users and user groups to applications.

This mapping controls which applications users or user groups can launch from a GlobalProtect Clientless VPN session.

The GlobalProtect portal uses the user/user group settings that you specify to determine which configuration to deliver to the GlobalProtect Clientless VPN user that connects. If you have multiple configurations, make sure they are ordered properly and map to all of the required applications, as the portal looks for a configuration match starting from the top of the list. As soon as the portal finds a match, it delivers the configuration to the GlobalProtect Clientless VPN user.

© 2017 Palo Alto Networks, Inc.



Publishing an application to a user/user group or allowing them to launch unpublished applications does not imply that they can access those applications. Controlling access to applications (published or not) is done using security policies.

You must configure group mapping (Device > User Identification > Group Mapping Settings) before you can select the groups.

1. In the **Applications** tab, **Add** an **Applications to User Mapping** to match users with published applications.
   - **Name**—Enter a name for the mapping (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
   - **Display application URL address bar**—Select this option to display an application URL address bar from which users can launch applications that are not published on the applications landing page. When enabled, users can click the **Application URL** link on the page and specify a URL.
2. Specify the **Source Users**. You can **Add** individual users or user groups to which the current application configuration applies. These users have permission to launch the configured applications using a GlobalProtect Clientless VPN. In addition to users and groups, you can specify when these settings apply to the users or groups:
   - **any**—The application configuration applies to all users (no need to **Add** users or user groups).
   - **select**—The application configuration applies only to users and user groups you **Add** to this list.
3. **Add** individual applications or application groups to the mapping. The **Source Users** you included in the configuration can use GlobalProtect Clientless VPN to link to the applications you add.

STEP 6 | Specify the security settings for a Clientless VPN session.

1. In the **Crypto Settings** tab, specify the authentication and encryption algorithms for the SSL sessions between the firewall and the published applications.
   - **Protocol Versions**—Select the required minimum and maximum TLS/SSL versions. The higher the TLS version, the more secure the connection. Choices include **SSLv3**, **TLSv1.0**, **TLSv1.1**, or **TLSv1.2**.

© 2017 Palo Alto Networks, Inc.

- **Key Exchange Algorithms**—Select the supported algorithm types for key exchange. Choices are: **RSA**, Diffie-Hellman (**DHE**), or Elliptic Curve Ephemeral Diffie-Hellman (**ECDHE**).
- **Encryption Algorithms**—Select the supported encryption algorithms. **AES128** or higher is recommended.
- **Authentication Algorithms**—Select the supported authentication algorithms. Choices are: **MD5**, **SHA1**, **SHA256**, or **SHA384**. **SHA256** or higher is recommended.

2. Select the action to take when the following issues occur with a server certificate presented by an application:
   - **Block sessions with expired certificate**—If the server certificate has expired, block access to the application.
   - **Block sessions with untrusted issuers**—If the server certificate is issued from an untrusted certificate authority, block access to the application.
   - **Block sessions with unknown certificate status**—If the OCSP or CRL service returns a certificate revocation status of `unknown`, block access to the application.
   - **Block sessions on certificate status check timeout**—If the certificate status check times out before receiving a response from any certificate status service, block access to the application.

STEP 7 | (**Optional**) Specify one or more proxy server configurations to access the applications.



Only basic authentication to the proxy is supported (username and password).

If users need to reach the applications through a proxy server, specify a **Proxy Server**. You can add multiple proxy server configurations, one for each set of domains.

- **Name**—A label (up to 31 characters) to identify the proxy server configuration. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
- **Domains**—Add the domains served by the proxy server. You can use a wild card character (*) at the beginning of the domain name to indicate multiple domains.
- **Use Proxy**—Select to assign a proxy server to provide access to the domains.
- **Server**—Specify the IP address or host name of the proxy server.
- **Port**—Specify a port for communication with the proxy server.
- **User** and **Password**—Specify the **User** and **Password** credentials needed to log in to the proxy server. Specify the password again for verification.

STEP 8 | (**Optional**) Specify any special treatment for application domains.

The Clientless VPN acts as a reverse proxy and modifies web pages returned by the published web applications. It rewrites all URLs and presents a rewritten page to remote users such that when they access any of those URLs, the requests go through GlobalProtect portal.

In some cases, the application may have pages that do not need to be accessed through the portal (for example, the application may include a stock ticker from yahoo.finance.com). You can exclude these pages.

In the **Advanced Settings** tab, **Add** domain names, host names, or IP addresses to the **Rewrite Exclude Domain List.** These domains are excluded from rewrite rules and cannot be rewritten.

Paths are not supported in host and domain names. The wildcard character (*) for host names and domain names can only appear at the beginning of the name (for example, *.etrade.com).

STEP 9 | Save the portal configuration.

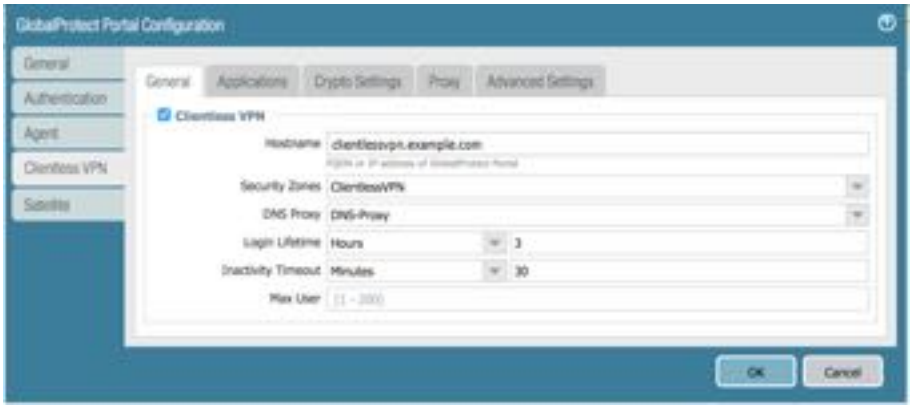
1. Click **OK** twice.
2. **Commit** your changes.

© 2017 Palo Alto Networks, Inc.

STEP 10 | Configure a Security policy rule to enable users to access the published applications.

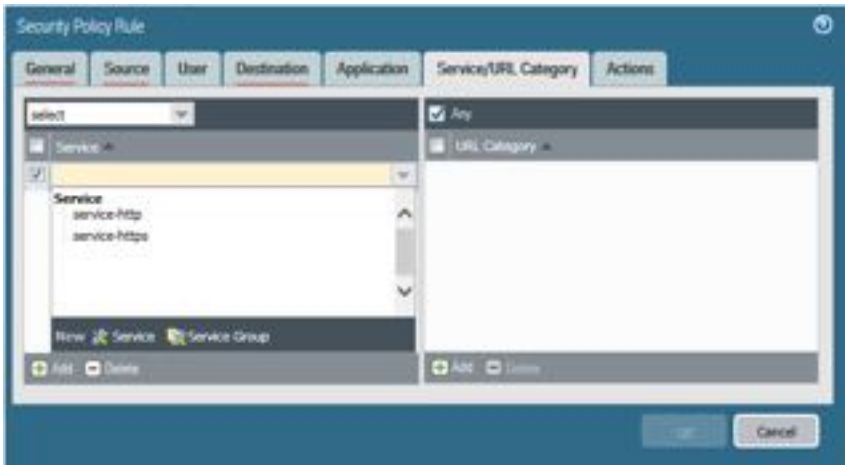
You need security policies for the following:

- Make the GlobalProtect portal which hosts Clientless VPN reachable from the Internet. This is traffic from Untrust or Internet Zone to the zone where you host the Clientless VPN portal.
- Allow Clientless VPN users to reach the Internet. This is traffic from the Clientless VPN zone to the Untrust or Internet Zone.



- Allow Clientless VPN users to reach corporate resources. This is traffic from the Clientless VPN zone to the Trust or Corp Zone. The security policies you define control which users have permission to use each published application. For the security zone where the published application servers are hosted, make sure **Enable User Identification** is set in order to create user-based rules for accessing published applications.

  By default **Service/URL** in **Security Policy Rule** is set **application-default**. Clientless VPN will not work for HTTPS sites with this default setting. Change **Service/URL** to include both **service-http** and **service-https**.



- When you configure a proxy server to access Clientless VPN applications, make sure you include the proxy IP address and port in the security policy definition. When applications are accessed through a proxy server, only security policies defined for the proxy IP address and port are applied.

© 2017 Palo Alto Networks, Inc.

Troubleshoot Clientless VPN

Because this feature involves dynamic re-writing of HTML applications, the HTML content for some applications may not re-write correctly and break the application. If issues occur, use the commands in the following table to help you identify the likely cause:

Table 4: Table: Rewrite Engine Statistics

Action	Command
CLI Commands	

List the version of Clientless VPN dynamic content being used
You can also view the dynamic update version from the **Device** > **Dynamic Updates** > **GlobalProtect Clientless VPN**.
	`pancpe@cagp> show system setting ssl-decrypt memory`
	`proxy uses shared allocator`
	`SSL certificate cache:`
	`Current Entries: 1`
	`Allocated 1, Freed 0`
	`Current CRE (61-62) : 3456 KB (Actual 3343 KB)`
	`Last CRE (60-47) : 3328 KB (Actual 3283 KB)`
	In this example, the current dynamic update is version 61-62, and the last installed dynamic update is version 60-47.

List active (current) users of Clientless VPN	`pancpe@cagp> show global-protect-portal current-user portal GPClientlessPortal filter-user all-users`
	`GlobalProtect Portal : GPClientlessPortal`
	`Vsys-Id : 1`
	`User : paloaltonetworks.com\johndoe`
	`Session-id : 1SU2vrPIDfdopGf-7gahMTCiX8PuL0S0`
	`Client-IP : 5.5.5.5`
	`Inactivity Timeout : 1800`
	`Seconds before inactivity timeout : 1750`
	`Login Lifetime : 10800`
	`Seconds before login lifetime : 10748`
	`Total number of user sessions: 1`

Show DNS resolution results
This can be useful to determine if there are DNS issues. If there is a DNS issue, you will notice querying against an FQDN that was not
	`pancpe@cagp> show system setting ssl-decrypt dns-cache`
	`Total DNS cache entries: 89`
	`Site IP Expire(secs) Interface`
	`bugzilla.panw.local 10.0.2.15 querying 0`
	`www.google.com 216.58.216.4 Expired 0`

© 2017 Palo Alto Networks, Inc.

Action	Command
resolvable in the CLI output.	`stats.g.doubleclick.net  74.125.199.154  Expired 0`

Show all Clientless VPN user sessions and cookies stored	`pancpe@cagp> show system setting ssl-decrypt gp-cookie-cache`
	
	`User: johndoe, Session-id: 1SU2vrPIDfdopGf-7gahMTCiX8PuL0S0, Client-ip: 199.167.55.50`

Show rewrite-stats
 	
This is useful to identify the health of the Clientless VPN rewrite engine.
 	
Refer to Table: Rewrite Engine Statistics for information on rewrite statistics and their meaning or purpose.
	`pancpe@cagp> show system setting ssl-decrypt rewrite-stats`
	
	`Rewrite Statistics`
	`initiate_connection : 11938`
	`setup_connection : 11909`
	`session_notify_mismatch : 1`
	`reuse_connection : 37`
	`file_end : 4719`
	`packet : 174257`
	`packet_mismatch_session : 1`
	`peer_queue_update_rcvd : 167305`
	`peer_queue_update_sent : 167305`
	`peer_queue_update_rcvd_failure: 66`
	`setup_connection_r : 11910`
	`packet_mismatch_session_r : 22`
	`pkt_no_dest : 23`
	`cookie_suspend : 2826`
	`cookie_resume : 2826`
	`decompress : 26`
	`decompress_freed : 26`
	`dns_resolve_timeout : 27`
	`stop_openend_response : 43`
	`received_fin_for_pending_req : 26`
	`Destination Statistics`
	`To mp : 4015`
	`To site : 12018`
	`To dp : 17276`
	`Return Codes Statistics`
	`ABORT : 18`
	`RESET : 30`
	`PROTOCOL_UNSUPPORTED : 7`
	`DEST_UNKNOWN : 10`
	`CODE_DONE : 52656`
	`DATA_GONE : 120359`
	`SWITCH_PARSER : 48`
	`INSERT_PARSER : 591`
	`SUSPEND : 2826`
	`Total Rewrite Bytes : 611111955`
	`Total Rewrite Useconds : 6902825`
	`Total Rewrite Calls : 176545`

Debug Commands	

© 2017 Palo Alto Networks, Inc.

Action	Command

Enable debug logs on the firewall running Clientless VPN Portal	`debug dataplane packet-diag set log feature ssl all`
	`debug dataplane packet-diag set log feature misc all`
	`debug dataplane packet-diag set log feature proxy all`
	`debug dataplane packet-diag set log feature flow basic`
	`debug dataplane packet-diag set log on`

Enable packet capture on the firewall running Clientless VPN Portal	`debug dataplane packet-diag set capture username <portal-username>`
	`debug dataplane packet-diag set capture stage clientless-vpn-client file clss_client1.pcap`
	`debug dataplane packet-diag set capture stage clientless-vpn-server file clss_server1.pcap`
	`debug dataplane packet-diag set capture stage firewall file clss_fw1.pcap`
	`debug dataplane packet-diag set capture stage receive file clss_rx1.pcap`
	`debug dataplane packet-diag set capture stage transmit file clss_tx1.pcap`
	`debug dataplane packet-diag set capture on`

Table 5: Table: Rewrite Engine Statistics

Statistic	Description
initiate_connection_failure	Connection initiation failed to back-end host
setup_connection_failure	Connection setup failed
setup_connection_duplicate	Duplicate peer session exists
session_notify_mismatch	Mostly invalid session
packet_mismatch_session	Failed to find right session for incoming packet
peer_queue_update_rcvd_failure	Session was invalid when packet update received by peer
peer_queue_update_sent_failure	Failed to send packet updates to peer or failed to send packet queue length updates to peer
exceed_pkt_queue_limit	Too many packets queued
proxy_connection_failure	Proxy connection failed
setup_connection_r	Installing the peer session to the application server. This value should match the values for **`initiate_connection`** and **`setup_connection`**.
setup_connection_duplicate_r	Duplicate sessions already in proxy

© 2017 Palo Alto Networks, Inc.

Statistic	Description
setup_connection_failure_r	Failed to set up the peer session
session_notify_mismatch_r	Peer session not found
packet_mismatch_session_r	Peer session not found when trying to get the packet
exceed_pkt_queue_limit_r	Too many packets held
unknown_dest	Failed to find destination host
pkt_no_dest	No destination for this packet
cookie_suspend	Suspended session to fetch cookies
cookie_resume	Received response from MP with updated cookies. This value generally matches the value of cookie_suspend.
decompress_failure	Failed to decompress
memory_alloc_failure	Failed to allocate memory
wait_for_dns_resolve	Suspended session to resolve DNS requests
dns_resolve_reschedule	Rescheduled DNS query due to no response (retry before timeout)
dns_resolve_timeout	DNS query timeout
setup_site_conn_failure	Failed to setup connection to site (proxy, DNS)
site_dns_invalid	DNS resolve failed
multiple_multipart	Multi-part content-type processed
site_from_referer	Received the back-end host from referrer. This can indicate failed rewrite links from flash or other content which Clientless VPN does not rewrite.
received_fin_for_pending_req	Received FIN from server for pending request from client
unmatched_http_state	Unexpected HTTP content. This can indicate an issue parsing the http headers or body.

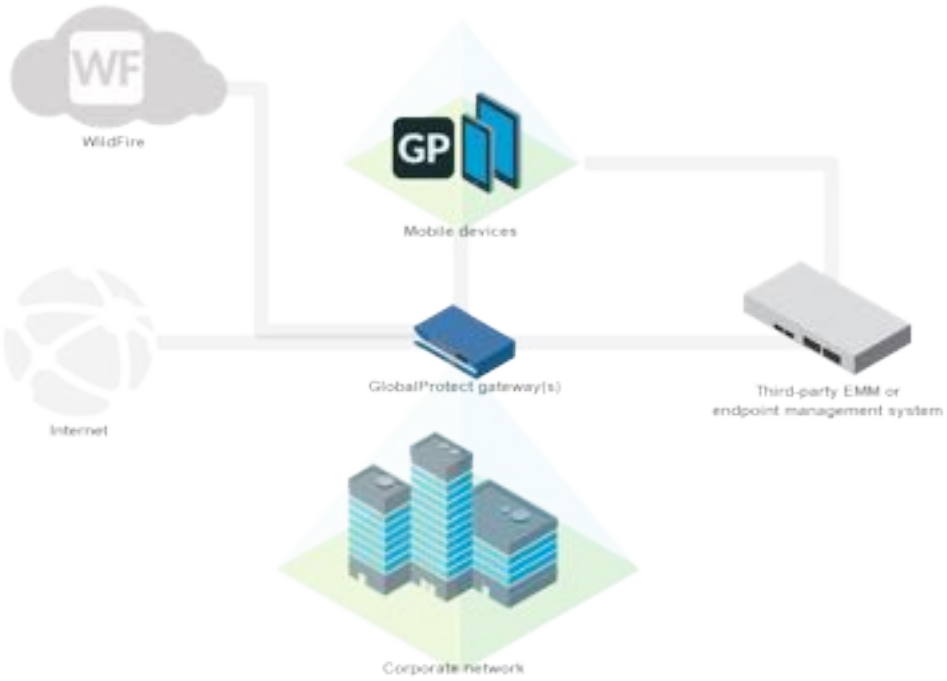
Mobile Endpoint Management

- Mobile Endpoint Management Overview on page 151
- Set Up a Mobile Endpoint Management System on page 152
- Deploy the GlobalProtect Mobile App Using AirWatch on page 153
- Manage the GlobalProtect App Using AirWatch on page 153
- Manage the GlobalProtect App Using a Third-Party MDM on page 164
- Configure Windows User-ID Agent to Collect Host Information on page 192

© 2017 Palo Alto Networks, Inc.

Mobile Endpoint Management Overview

As mobile endpoints become more powerful, end users increasingly rely on them to perform business tasks. However, these same endpoints that access your corporate network also connect to the internet without protection against threats and vulnerabilities. By using a third-party mobile endpoint management system—such as a mobile device management (MDM) or enterprise mobility management (EMM) system—you can easily manage both company-provisioned and employee-owned devices (such as in a BYOD environment).



A mobile endpoint management system simplifies the administration of mobile endpoints by enabling you to automatically deploy your corporate account configuration and VPN settings to compliant endpoints. You can also use your mobile endpoint management system for remediation of security breaches by interacting with an endpoint that has been compromised. This protects both corporate data as well as personal end user data. For example, if an end user loses an endpoint, you can remotely lock the endpoint from the mobile endpoint management system or even wipe the endpoint (either completely or selectively).

In addition to the account provisioning and remote device management functions that a mobile endpoint management system can provide, when integrated with your existing GlobalProtect™ VPN infrastructure, you use host information that the endpoint reports to enforce security policies for access to apps through the GlobalProtect gateway. You can also use the monitoring tools that are built into the Palo Alto next-generation firewall to monitor mobile endpoint traffic.

© 2017 Palo Alto Networks, Inc.

Set Up a Mobile Endpoint Management System

To set up a mobile endpoint management system, use the following workflow:

STEP 1 | Set up the GlobalProtect Infrastructure.

1. Create Interfaces and Zones for GlobalProtect on page 15.
2. Enable SSL Between GlobalProtect Components on page 17.
3. Set up GlobalProtect User Authentication. Refer to About GlobalProtect User Authentication on page 27.
4. Enable Group Mapping on page 63.
5. Configure a GlobalProtect Gateway on page 72.
6. Activate Licenses for each firewall running a gateway(s) that supports the GlobalProtect app on mobile endpoints.
7. Set Up Access to the GlobalProtect Portal on page 83.

STEP 2 | Set up the mobile endpoint management system and decide whether to support only corporate-issued endpoints or both corporate-issued and personal endpoints.

See the instructions for your mobile endpoint management system, mobile device management (MDM) system, or enterprise mobility management (EMM) system.

STEP 3 | Obtain the GlobalProtect app for mobile endpoints.

- App store—Download and Install the GlobalProtect Mobile App on page 109
- AirWatch—Deploy the GlobalProtect Mobile App Using AirWatch on page 153
- Other third-party mobile endpoint management system—See the instructions from your vendor on how to deploy apps to managed endpoints.

STEP 4 | Configure VPN settings for the GlobalProtect app.

- Manage the GlobalProtect App Using AirWatch on page 153
- Manage the GlobalProtect App Using a Third-Party MDM on page 164

STEP 5 | Configure polices that target mobile endpoints using host information.

Configure HIP-Based Policy Enforcement on page 177 for managed endpoints.

© 2017 Palo Alto Networks, Inc.

Manage the GlobalProtect App Using AirWatch

- Deploy the GlobalProtect Mobile App Using AirWatch on page 153
- Configure the GlobalProtect App for iOS Using AirWatch on page 154
- Configure the GlobalProtect App for Android Using AirWatch on page 157
- Configure the GlobalProtect App for Windows 10 UWP Using AirWatch on page 161

Deploy the GlobalProtect Mobile App Using AirWatch

The GlobalProtect app provides a simple way to extend the enterprise security policies out to mobile endpoints. As with other remote hosts running the GlobalProtect agent, the mobile app provides secure access to your corporate network over an IPSec or SSL VPN tunnel. The app connects to the gateway that is closest to the end user's current location. In addition, traffic to and from the mobile endpoint is automatically subject to the same security policy enforcement as other hosts on your corporate network. Like the GlobalProtect agent, the app collects information about the host configuration and can use this information for enhanced HIP-based security policy enforcement.

There are two primary methods for installing the GlobalProtect app: You can you can install the app directly from the app store for your endpoint (see Download and Install the GlobalProtect Mobile App on page 109); or, deploy the app from a third-party mobile endpoint management system (such as AirWatch) and transparently push the app to your managed endpoints.

With AirWatch, you can deploy the GlobalProtect app to managed endpoints that have enrolled with AirWatch. Endpoints running iOS or Android must download the AirWatch agent to enroll with the AirWatch EDM. Windows 10 endpoints do not require the AirWatch agent but require you to configure enrollment on the endpoint. After you deploy the app, configure and deploy a VPN profile to set up the GlobalProtect app for the end user automatically.

STEP 1 | Before you begin, ensure that the endpoints to which you want to deploy the GlobalProtect app are enrolled with AirWatch:

- **Android and iOS**—Download the AirWatch agent and following the prompts to enroll.
- **Windows Phone and Windows 10 UWP**—Configure the Windows 10 UWP endpoint to enroll with AirWatch (from the endpoint, select **Settings** > **Accounts** > **Work access** > **Connect**).

STEP 2 | From AirWatch, select **Apps & Books** > **Public** > **Add Application**.

STEP 3 | Select the organization group by which this app will be managed.

STEP 4 | Select the **Platform**, either **Apple iOS**, **Android**, or **Windows Phone**.

STEP 5 | Search for the app in the app store for the endpoint or enter the URL of the GlobalProtect app page:

- **Apple iOS**—https://itunes.apple.com/us/app/globalprotect/id592489989?mt=8&uo=4
- **Android**—https://play.google.com/store/apps/details?id=com.paloaltonetworks.globalprotect
- **Windows Phone**—https://www.microsoft.com/store/apps/9NBLGGH6BZL3

STEP 6 | Click **Next**. If you chose to search for the app the app store for the endpoint, you must also **Select** the app from a list of search results.

© 2017 Palo Alto Networks, Inc.

If you chose to search for the GlobalProtect app for Android and did not see the app in the list, contact your Android for Work administrator to add GlobalProtect to the list of approved company apps.

STEP 7 | On the **Assignment** tab, select **Assigned Smart Groups** that will have access to this app.

STEP 8 | On the **Deployment** tab, select the **Push Mode**, either **Auto** or **On Demand**.

STEP 9 | Select **Save & Publish** to push the App Catalog to the endpoints in the Smart Groups you assigned in the **Assignment** section.

STEP 10 | Next steps:

- Configure the GlobalProtect App for iOS Using AirWatch on page 154
- Configure the GlobalProtect App for Android Using AirWatch on page 157
- Configure the GlobalProtect App for Windows 10 UWP Using AirWatch on page 161

Configure the GlobalProtect App for iOS Using AirWatch

AirWatch is an Enterprise Mobility Management Platform that enables you to manage mobile endpoints, from a central console. The GlobalProtect app provides a secure connection between AirWatch managed mobile endpoints and the firewall at either the device or application level. Using GlobalProtect as the secure connection allows consistent inspection of traffic and enforcement of network security policy for threat prevention on the mobile endpoint.

- Configure a Device-Level VPN Configuration for iOS Devices Using AirWatch on page 154
- Configure a Per-App VPN Configuration for iOS Devices Using AirWatch on page 155

Configure a Device-Level VPN Configuration for iOS Devices Using AirWatch

You can easily enable access to internal resources from your managed mobile endpoints by configuring VPN access using AirWatch. In a device-level VPN configuration, you route all of the traffic that matches the access routes configured on the GlobalProtect gateway through the GlobalProtect VPN.

STEP 1 | Download the GlobalProtect app for iOS.

- Deploy the GlobalProtect Mobile App Using AirWatch on page 153.
- Download the GlobalProtect app directly from the App Store.

STEP 2 | From the AirWatch console, modify or add a new Apple iOS profile.

1. Navigate to **Devices** > **Profiles** > **List View**.
2. Select an existing profile to add the VPN configuration to it or add a new one (select **Add** > **Apple iOS)**.
3. Configure **General** profile settings:

- **Description**—A brief description of the profile that indicates its purpose.
- **Deployment**—Determines if the profile will be automatically removed upon unenrollment, either **Managed** (the profile is removed) or **Manual** (the profile remains installed until removed by the end user).
- **Assignment Type**—Determines how the profile is deployed to endpoints. Select **Auto** to deploy the profile to all endpoints automatically, **Optional** to enable the end user to install the profile from the Self-Service Portal (SSP) or to manually deploy the profile to individual endpoints, or **Compliance** to deploy the profile when an end user violates a compliance policy applicable to the endpoint.

© 2017 Palo Alto Networks, Inc.

- **Managed By**—The Organization Group with administrative access to the profile.
- **Assigned Smart Group**—The Smart Group to which you want the device profile added. Includes an option to create a new Smart Group which can be configured with specs for minimum OS, device models, ownership categories, organization groups and more.
- **Allow Removal**—Determines whether or not the profile can be removed by the endpoint's end user. Select **Always** to enable the end user to manually remove the profile at any time, **Never** to prevent the end user from removing the profile from the endpoint, or **With Authorization** to enable the end user to remove the profile with the authorization of the administrator. Choosing **With Authorization** adds a required Password.
- **Exclusions**—If **Yes** is selected, a new field **Excluded Smart Groups** displays, enabling you to select those Smart Groups you wish to exclude from the assignment of this device profile.

STEP 3 | To configure the VPN settings, select **VPN** and then click **Configure**.

STEP 4 | Configure connection information, including:

- **Connection Name**—Enter the name of the connection name to be displayed.
- **Connection Type**—Select **Palo Alto Networks GlobalProtect** as the network connection method.
- **Server**—Enter the hostname or IP address of the GlobalProtect portal to which to connect.
- **Account**—Enter the username of the VPN account or click add ("**+**") to view supported lookup values you can insert.
- **Authentication**—Choose the method to authenticate end users. Follow the related prompts to enter a **Password** or upload an **Identity Certificate** to use to authenticate users; Or, if you selected **Password + Certificate**, follow the related prompts for both.

STEP 5 | **Save & Publish** your changes.

Configure a Per-App VPN Configuration for iOS Devices Using AirWatch

You can easily enable access to internal resources from your managed mobile endpoints by configuring GlobalProtect VPN access using AirWatch. In a per-app VPN configuration, you can specify which managed apps on the endpoint can send traffic through the GlobalProtect VPN tunnel. Unmanaged apps will continue to connect directly to the Internet instead of through the GlobalProtect VPN tunnel.

STEP 1 | Download the GlobalProtect app for iOS:

- Deploy the GlobalProtect Mobile App Using AirWatch on page 153.
- Download the GlobalProtect app directly from the App Store.

STEP 2 | From the AirWatch console, modify or add a new Apple iOS profile:

1. Navigate to **Devices** > **Profiles** > **List View**.
2. Select an existing profile to add the VPN configuration to it or add a new one (select **Add** > **Apple iOS**).

STEP 3 | Configure **General** profile settings:

- **Description**—A brief description of the profile that indicates its purpose.
- **Deployment**—Determines if the profile will be automatically removed upon unenrollment, either **Managed** (the profile is removed) or **Manual** (the profile remains installed until removed by the end user).
- **Assignment Type**—Determines how the profile is deployed to endpoints. Select **Auto** to deploy the profile to all endpoints automatically, **Optional** to enable the end user to install the profile from the Self-Service Portal (SSP) or to manually deploy the profile to individual endpoints, or **Compliance** to deploy the profile when an end user violates a compliance policy applicable to the endpoint.

© 2017 Palo Alto Networks, Inc.

- **Managed By**—The Organization Group with administrative access to the profile.
- **Assigned Smart Group**—The Smart Group to which you want the device profile added. Includes an option to create a new Smart Group which can be configured with specs for minimum OS, device models, ownership categories, organization groups and more.
- **Allow Removal**—Determines whether or not the profile can be removed by the endpoint's end user. Select **Always** to enable the end user to manually remove the profile at any time, **Never** to prevent the end user from removing the profile from the endpoint, or **With Authorization** to enable the end user to remove the profile with the authorization of the administrator. Choosing **With Authorization** adds a required Password.
- **Exclusions**—If **Yes** is selected, a new field **Excluded Smart Groups** displays, enabling you to select those Smart Groups you wish to exclude from the assignment of this device profile.

STEP 4 | To configure the per-app VPN settings in the Apple iOS profile, select **VPN** and then click **Configure**.

STEP 5 | Configure connection information, including:

- **Connection Name**—Enter the name of the connection name to be displayed.
- **Connection Type**—Select **Palo Alto Networks GlobalProtect** as the network connection method.
- **Server**—Enter the hostname or IP address of the GlobalProtect portal to which to connect.
- **Account**—Enter the username of the VPN account or click add ("**+**") to view supported lookup values that you can insert.
- **Send All Traffic**—Select this check box to force all traffic through the specified network.
- **Disconnect on Idle**—Allow the VPN to auto-disconnect after a specific amount of time.
- Enable **Per App VPN** to route all of the traffic for a managed app traffic through the GlobalProtect VPN.
- **Connect Automatically**—Select this check box to allow the VPN to connect automatically to chosen Safari Domains.

STEP 6 | Configure the authentication method to use to authenticate users. For per-app VPN, you must use certificate-based authentication. Select **User Authentication: Certificate**, and then follow the prompts to upload an **Identity Certificate** to use for authentication.

STEP 7 | Select either **Manual** or **Auto Proxy** type and enter the specific information needed.

STEP 8 | Click **Save & Publish**.

STEP 9 | Configure per-app VPN settings for a new managed app, or modify the settings for an existing managed apps.

After configuring the settings for the app and enabling per-app VPN, you can publish the app to a group of users and enable the app to send traffic through the GlobalProtect VPN tunnel.

1. On the main page, select **Apps & Books** > **Public.**
2. To add a new app, select **Add Application**. Or, to modify the settings of an existing app, locate the GlobalProtect app in the list of Public apps and then select the edit icon in the actions menu next to the row.
3. Select the organization group by which this app will be managed.
4. Select **Apple iOS** as the **Platform**.
5. Select your preferred method for locating the app, either by searching the App Store (by Name), or specifying a URL for the app in the App Store (for example, to add the Box app, enter https://itunes.apple.com/us/app/box-for-iphone-and-ipad/id290853822?mt=8&uo=4), and then click **Next**. If you choose to search the App Store, you must **Select** the app from the list of search results.
6. On the **Assignment** tab, select **Assigned Smart Groups** that will have access to this app.

© 2017 Palo Alto Networks, Inc.

7. On the **Deployment** tab, select the **Push Mode**, either **Auto** or **On Demand**.
8. Select **Use VPN** and then select the Apple iOS profile that you created earlier in this workflow.



Only profiles that have per-app VPN enabled are available from the drop-down.

1. Select **Save & Publish** to push the App Catalog to the endpoints in the Smart Groups you assigned in the **Assignment** section.

Configure the GlobalProtect App for Android Using AirWatch

You can use the GlobalProtect App for Android with AirWatch agent 6.0 and later releases. The AirWatch agent interfaces with AirWatch to manage Android endpoints. Using the GlobalProtect app for Android as the secure connection between the endpoint and the firewall allows consistent inspection of traffic and enforcement of network security policy for threat prevention. The GlobalProtect app can provide a secure connection at either the device or application level.

- Configure a Device-Level VPN Configuration for Android Devices Using AirWatch on page 157
- Configure a Per-App VPN Configuration for Android Devices Using AirWatch on page 158
- Enable App Scan Integration with WildFire on page 160

Configure a Device-Level VPN Configuration for Android Devices Using AirWatch

You can easily enable access to internal resources from your managed Android mobile endpoints by configuring VPN access using AirWatch. In a device-level VPN configuration, you route all of the traffic that matches the access routes configured on the GlobalProtect gateway through the GlobalProtect VPN.

STEP 1 | Download the GlobalProtect app for Android:

- Deploy the GlobalProtect Mobile App Using AirWatch on page 153.
- Download the GlobalProtect app directly from Google Play.

STEP 2 | From the AirWatch console, modify or add a new Android profile.

1. Navigate to **Devices** > **Profiles** > **List View**.
2. Select an existing profile to which to add the VPN configuration or add a new one (select **Add** > **Add Profile**).
3. Select **Android** as the platform and **Device** as the configuration type.

STEP 3 | Configure **General** profile settings:

- **Name**—Provide a meaningful name for this configuration.
- **Version**—This field is auto-populated with the latest version number of the configuration profile.
- **Description**—A brief description of the profile that indicates its purpose.
- Profile Scope—Scope for this profile, either **Production, Staging**, or **Both**.
- **Assignment Type**—Determines how the profile is deployed to endpoints. Select **Auto** to deploy the profile to all endpoints automatically, **Optional** to enable the end user to install the profile from the Self-Service Portal (SSP) or to manually deploy the profile to individual endpoints, or **Compliance** to deploy the profile when an end user violates a compliance policy applicable to the endpoint.
- **Managed By**—The Organization Group with administrative access to the profile.
- **Assigned Smart Group**—The Smart Group to which you want the device profile added. Includes an option to create a new Smart Group which can be configured with specs for minimum OS, device models, ownership categories, organization groups and more.

© 2017 Palo Alto Networks, Inc.

- **Allow Removal**—Determines whether or not the profile can be removed by the endpoint's end user. Select **Always** to enable the end user to manually remove the profile at any time, **Never** to prevent the end user from removing the profile from the endpoint, or **With Authorization** to enable the end user to remove the profile with the authorization of the administrator. Choosing **With Authorization** adds a required Password.
- **Exclusions**—If **Yes** is selected, a new field **Excluded Smart Groups** displays, enabling you to select those Smart Groups you wish to exclude from the assignment of this device profile.

STEP 4 | **Save and Publish** this profile to the assigned Smart Groups.

STEP 5 | To configure the VPN settings, select **VPN** and then click **Configure**.

STEP 6 | Configure **Connection Info**, including:

- **Connection Type**—Select **GlobalProtect** as the network connection method.
- **Connection Name**—Enter the name of the connection name that the endpoint will display.
- **Server**—Enter the hostname or IP address of the GlobalProtect portal to which to connect.

STEP 7 | Configure **Authentication** information:

1. Choose the method to authenticate end users: **Password** or **Certificate**.
2. Enter the **Username** of the VPN account or click add ("**+**") to view supported lookup values that you can insert.
3. Enter a **Password** or upload an **Identity Certificate** that GlobalProtect will use to authenticate users.

STEP 8 | **Save & Publish** this profile to the assigned Smart Groups.

Configure a Per-App VPN Configuration for Android Devices Using AirWatch

You can easily enable access to internal resources from your managed mobile endpoints by configuring GlobalProtect VPN access using AirWatch. In a per-app VPN configuration, you can specify which managed apps on the endpoint can send traffic through the GlobalProtect VPN tunnel. Unmanaged apps will continue to connect directly to the Internet instead of through the GlobalProtect VPN tunnel.

STEP 1 | Download the GlobalProtect app for Android:

- Deploy the GlobalProtect Mobile App Using AirWatch on page 153.
- Download the GlobalProtect app directly from Google Play.

STEP 2 | From the AirWatch console, modify or add a new Android profile.

1. Navigate to **Devices** > **Profiles** > **List View**.
2. Select an existing profile to which to add the VPN configuration or add a new one (select **Add** > **Add Profile)**.
3. Select **Android** as the platform and **Device** as the configuration type.

STEP 3 | Configure **General** profile settings:

- **Name**—Provide a meaningful name for this configuration.
- **Version**—This field is auto-populated with the latest version number of the configuration profile.
- **Description**—A brief description of the profile that indicates its purpose.
- Profile Scope—Scope for this profile, either **Production, Staging**, or **Both**.
- **Assignment Type**—Determines how the profile is deployed to endpoints. Select **Auto** to deploy the profile to all endpoints automatically, **Optional** to enable the end user to install the profile from the

© 2017 Palo Alto Networks, Inc.

Self-Service Portal (SSP) or to manually deploy the profile to individual endpoints, or **Compliance** to deploy the profile when an end user violates a compliance policy applicable to the endpoint.

- **Managed By**—The Organization Group with administrative access to the profile.
- **Assigned Smart Group**—The Smart Group to which you want the device profile added. Includes an option to create a new Smart Group which can be configured with specs for minimum OS, device models, ownership categories, organization groups and more.
- **Allow Removal**—Determines whether or not the profile can be removed by the endpoint's end user. Select **Always** to enable the end user to manually remove the profile at any time, **Never** to prevent the end user from removing the profile from the endpoint, or **With Authorization** to enable the end user to remove the profile with the authorization of the administrator. Choosing **With Authorization** adds a required **Password**.
- **Exclusions**—When you select **Yes**, the AirWatch console displays an **Excluded Smart Groups** field which you can use to select those Smart Groups you wish to exclude from the assignment of this device profile.

STEP 4 | **Save and Publish** this profile to the assigned Smart Groups.

STEP 5 | To configure the VPN settings:

1. Select **VPN** and then click **Configure**.
2. Configure **Connection Info**, including:

- **Connection Type**—Select **GlobalProtect** as the network connection method.
- **Connection Name**—Enter the name of the connection name that the endpoint will display.
- **Server**—Enter the hostname or IP address of the GlobalProtect portal to which to connect.
- Enable **Per App VPN** to route all of the traffic for a managed app traffic through the GlobalProtect VPN.

1. Select the authentication method to use to authenticate users. For per-app VPN, you must use certificate-based authentication. Select **User Authentication: Certificate**, and then follow the prompts to upload an **Identity Certificate** to use for authentication.
2. **Save & Publish** this profile to the assigned Smart Groups.

STEP 6 | Configure per-app VPN settings for a new managed app, or modify the settings for an existing managed apps:

1. On the main page, select **Apps & Books** > **Applications** > **List View** > **Public**.
2. To add a new app, select **Add Application**. Or, to modify the settings of an existing app, locate the app in the list of Public apps and then select the edit icon in the actions menu next to the row.
3. Select the organization group by which this app will be managed.
4. Select **Android** as the **Platform**.
5. Select your preferred method for locating the app, either by specifying a URL or importing the app from the app store (Google Play). To search by URL, you must also enter the Google Play Store URL for the app (for example, to search for the Box app by URL, enter https://play.google.com/store/apps/details?id=com.box.android).
6. Click **Next**. If you chose to import the app from Google Play in the previous step, you must **Select** the app from the list of approved company apps. If you do not see the app in the list, contact your Android for Work administrator to approve the app.
7. On the **Assignment** tab, select **Assigned Smart Groups** that will have access to this app.
8. On the **Deployment** tab, select the **Push Mode**, either **Auto** or **On Demand**.
9. Select **Use VPN** and then select the Android profile that you created earlier in this workflow.

Only profiles that have per-app VPN enabled are available from the drop-down.

© 2017 Palo Alto Networks, Inc.

1. **Save & Publish** the configuration to the assigned Smart Groups.

STEP 7 | Configure **Authentication** information:

1. Choose the method to authenticate end users: **Password** or **Certificate**.
2. Enter the **Username** of the VPN account or click add (“**+**”) to view supported lookup values that you can insert.
3. Enter a **Password** or upload an **Identity Certificate** that GlobalProtect will use to authenticate users.

STEP 8 | **Save & Publish** this profile to the assigned Smart Groups.

Enable App Scan Integration with WildFire

By enabling App Scan in AirWatch, you can leverage WildFire threat intelligence about apps to detect malware on Android endpoints. When enabled, the AirWatch agent sends the list of apps that are installed on the Android endpoint to AirWatch. This occurs during enrollment and subsequently on any device check-in. AirWatch then periodically queries WildFire for verdicts and can take compliance action on the endpoint based on the verdict.

STEP 1 | Before you begin, obtain a WildFire API key. If you do not already have an API key, contact Support.

STEP 2 | From AirWatch, select **Groups & Settings** > **All Settings** > **Apps** > **App Scan** > **Third Party Integration**.

STEP 3 | Select **Current Setting: Override**.

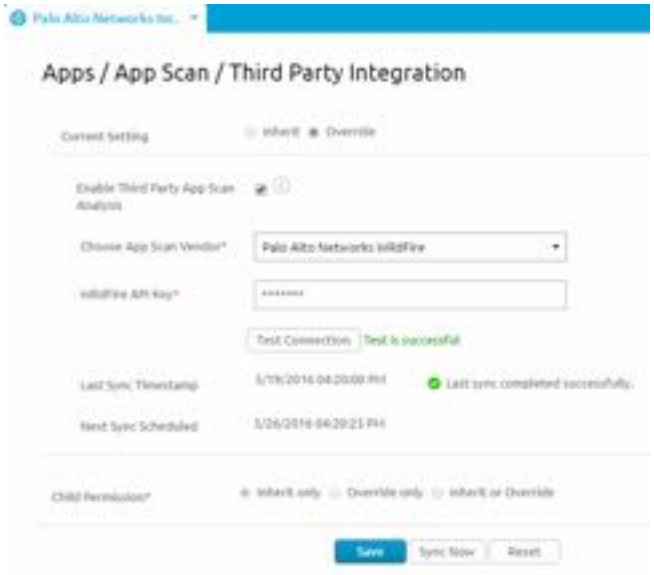
STEP 4 | Select **Enable Third Party App Scan Analysis** to enable communication between AirWatch and WildFire.

STEP 5 | Choose **Palo Alto Networks WildFire** from the **Choose App Scan Vendor** drop-down.

STEP 6 | Enter your WildFire API key.

STEP 7 | Click **Test Connection** to ensure that AirWatch can communicate with WildFire. If the test is not successful, verify connectivity to the Internet, re-enter the API key, and then try again.

© 2017 Palo Alto Networks, Inc.



STEP 8 | **Save** your changes. AirWatch schedules a synchronization task to communicate with WildFire to obtain the latest verdicts for application hashes and runs the task at regular intervals. Click **Sync Now** to initiate a manual sync with WildFire.

Configure the GlobalProtect App for Windows 10 UWP Using AirWatch

Using the GlobalProtect app for Windows10UWP as the secure connection between the endpoint and the firewall allows consistent inspection of traffic and enforcement of network security policy for threat prevention.

The GlobalProtect app for Windows 10 UWP supports the following configurations using AirWatch:

- **Per-App VPN**—Specifies which managed apps on the endpoint can send traffic through the secure tunnel. Unmanaged apps will continue to connect directly to the Internet instead of through the secure connection.
- **Device-Level VPN**—Sends all traffic that matches specific filters (such as port and IP address) through the VPN irrespective of app. Device-level VPN configurations also support the ability to force the secure connection to be **Always On**. For even tighter security requirements, you can enable the **VPN Lockdown** option which both forces the secure connection to always be on and connected and disables network access when the app is not connected. This configuration is similar to the **Enforce GlobalProtect for Network Access** option that you would typically configure in a GlobalProtect portal configuration.

Because AirWatch does not yet list GlobalProtect as an official connection provider for Windows endpoints, you must select an alternate VPN provider, edit the settings for the GlobalProtect app, and import the configuration back into the VPN profile as described in the following workflow.

STEP 1 | Download the GlobalProtect app for Windows 10 UWP:

- Deploy the GlobalProtect Mobile App Using AirWatch.
- Download the GlobalProtect app directly from the Microsoft Store.

STEP 2 | From the AirWatch console, add a new Windows 10 UWP profile:

© 2017 Palo Alto Networks, Inc.

1. Navigate to **Devices** > **Profiles** > **List View**.
2. Select **Add** > **Add Profile**.
3. Select **Windows** as the platform and **Windows Phone** as the configuration type.
4. Configure **General** profile settings such as a meaningful **Name** for this configuration and a brief **Description** of the profile that indicates its purpose.
5. **Save and Publish** this profile to the assigned Smart Groups.

STEP 3 | To configure the VPN connection settings, select **VPN** and then click **Configure**.

STEP 4 | Select Configure **Connection Info**, including:

- **Connection Name**—Enter the name of the connection name that the endpoint will display.
- **Connection Type**—Select an alternate provider (do not select **IKEv2**, **L2TP**, **PPTP**, or **Automatic** as these do not have the associated vendor settings required for the GlobalProtect VPN profile).

  You must select the alternate vendor because AirWatch does not yet list GlobalProtect as an official connection provider for Windows endpoints.
- **Server**—Enter the hostname or IP address of the GlobalProtect portal to which to connect.

STEP 5 | Configure the authentication settings for the VPN connection:

1. Select the **Authentication Type** to choose the method to authenticate end users.
2. To permit GlobalProtect to save user credentials, enable **Remember Credentials** in the Policies area.

STEP 6 | Configure VPN traffic rules to apply device wide or on a per-app basis:

- **Add New Per-App VPN Rule**—Specify rules for specific legacy apps (typically .exe files) or modern apps (typically downloaded from the Microsoft Store) that determine whether to automatically establish the VPN connection when the app is launched and whether to send app traffic through the VPN. You can also configure specific traffic filters to route only app traffic through the VPN if it matches match criteria such as IP address and port.
- **Add New Device-Wide VPN Rule**—Specify routing filters to send traffic matching a specific route through the VPN. These rules are not bound by application and are evaluated across the endpoint. If the traffic matches the match criteria, it is routed through the VPN.

STEP 7 | (**Device-level VPN only**) If desired, configure your preference of Always-On connection:

1. To maintain the VPN connection always, enable either of the following options:
   - **Always On**—Force the secure connection to be always on.
   - **VPN Lockdown**—Force the secure connection to be always on and connected, and disable the network access when the app is not connected. The **VPN Lockdown** option in AirWatch is similar to the **Enforce GlobalProtect for Network Access** option that you would configure in a GlobalProtect portal configuration.
2. Specify **Trusted Network** addresses if you want GlobalProtect to connect only when it detects a trusted network connection.
3. **Save & Publish** your changes.

STEP 8 | To adapt the configuration for GlobalProtect, edit the VPN profile in XML.

To minimize additional edits in the raw XML, review the settings in your VPN profile before you export the configuration. If you need to change a setting after you export the VPN profile, you can make the changes in the raw XML or, you can update the setting in the VPN profile and perform this step again.

© 2017 Palo Alto Networks, Inc.

1. In the **Devices** > **Profiles** > **List View**, select the radio button next to the new profile you added in the previous steps, and then select **</>XML** at the top of the table. AirWatch opens the XML view of the profile.
2. **Export** the profile and then open it in a text editor of your choice.
3. Edit the following settings for GlobalProtect:

- In the `LoclURI` element that specifies the `PluginPackageFamilyName`, change the element to:

  <LocURI>./Vendor/MSFT/VPNv2/PaloAltoNetworks/PluginProfile/PluginPackageFamilyName</LocURI>
- In the `Data` element that follows, change the value to:

  <Data>PaloAltoNetworks.GlobalProtect_rn9aeerfb38dg</Data>

1. Save your changes to the exported profile.
2. Return to AirWatch and the **Devices** > **Profiles** > **List View**.
3. Create (select **Add** > **Add Profile** > **Windows** > **Windows Phone**) and name a new profile.
4. Select **Custom Settings** > **Configure**, and then copy and paste the edited configuration.
5. **Save & Publish** your changes.

STEP 9 | Clean up the original profile: Select the original profile from the **Devices** > **Profiles** > **List View**, select **More Actions** > **Deactivate**. AirWatch moves the profile to the Inactive list.

STEP 10 | Test the configuration.

© 2017 Palo Alto Networks, Inc.

Manage the GlobalProtect App Using a Third-Party MDM

You can use any third-party mobile device management system, such as a mobile device management (MDM) system, that manages an Android or iOS mobile endpoint to deploy and configure the GlobalProtect app.

- Manage the GlobalProtect App for iOS Using a Third-Party MDM System
  - Configure the GlobalProtect App for iOS on page 164
  - Example: GlobalProtect iOS App Device-Level VPN Configuration on page 165
  - Example: GlobalProtect iOS App App-Level VPN Configuration on page 166
- Manage the GlobalProtect App for Android Using a Third-Party MDM System
  - Configure the GlobalProtect App for Android on page 168
  - Example: Set VPN Configuration on page 169
  - Example: Remove VPN Configuration on page 169

Configure the GlobalProtect App for iOS

While a third-party MDM system allows you to push configuration settings that allow access to your corporate resources and provides a mechanism for enforcing device restrictions, it does not secure the connection between the mobile endpoint and services it connects to. To enable the client to establish secure tunnel connections, you must enable VPN support on the endpoint.

The following table describes typical settings that you can configure using your third-party MDM system.

Setting	Description	Value
Connection Type	Type of connection enabled by the policy.	**`Custom SSL`**
Identifier	Identifier for the custom SSL VPN in reverse DNS format.	**`com.paloaltonetworks.GlobalProtect.vpnplugin`**
Server	Host name or IP address of the GlobalProtect portal.	**`<hostname or IP address>`** For example: **`gp.paloaltonetworks.com`**
Account	User account for authenticating the connection.	**`<username>`**
User Authentication	Authentication type for the connection.	**`Certificate`** | **`Password`**
Credential	(**Certificate User Authentication only**) Credential for authenticating the connection.	**`<credential>`** For example: **`clientcredial.p12`**

© 2017 Palo Alto Networks, Inc.

Setting	Description	Value
Enable VPN On Demand	(**Optional**) Domain and hostname that will establish the connection and the on-demand action: • Always establish a connection • Never establish a connection • Establish a connection if needed	**<domain and hostname and the on-demand action>** For example: `gp.acme.com; Never establish`

Example: GlobalProtect iOS App Device-Level VPN Configuration

The following example shows the XML configuration containing a VPN payload that you can use to verify the device-level VPN configuration of the GlobalProtect app for iOS.

```
<?xml version="1.0"
encoding="UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN" "http://www.apple.com/
DTDs/PropertyList-1.0.dtd">
<plist version="1.0">
<dict>
<key>PayloadContent</key>
<array>
<dict>
<key>PayloadDescription</key>
<string>Configures VPN settings, including authentication.</string>
<key>PayloadDisplayName</key>
<string>VPN (Sample Device Level VPN)</string>
<key>PayloadIdentifier</key>
<string>Sample Device Level VPN.vpn</string>
<key>PayloadOrganization</key>
<string>Palo Alto Networks</string>
<key>PayloadType</key>
<string>com.apple.vpn.managed</string>
<key>PayloadVersion</key>
<integer>1</integer>
<key>PayloadUUID</key>
<string>5436fc94-205f-7c59-0000-011d</string>
<key>UserDefinedName</key>
<string>Sample Device Level VPN</string>
<key>Proxies</key>
<dict/>
<key>VPNType</key>
<string>VPN</string>
<key>VPNSubType</key>
<string>com.paloaltonetworks.GlobalProtect.vpnplugin</string>
<key>IPv4</key>
<dict>
<key>OverridePrimary</key>
<integer>0</integer>
</dict>
<key>VPN</key>
<dict>
<key>RemoteAddress</key>
<string>cademogp.paloaltonetworks.com</string>
<key>AuthName</key>
<string></string>
<key>DisconnectOnIdle</key>
```


© 2017 Palo Alto Networks, Inc.

```
<integer>0</integer>
<key>OnDemandEnabled</key>
<integer>1</integer>
<key>OnDemandRules</key>
<array>
<dict>
<key>Action</key>
<string>Connect</string>
</dict>
</array>
<key>AuthenticationMethod</key>
<string>Password</string>
</dict>
<key>VendorConfig</key>
<dict>
<key>AllowPortalProfile</key>
<integer>0</integer>
<key>FromAspen</key>
<integer>1</integer>
</dict>
</dict>
</array>
<key>PayloadDisplayName</key>
<string>Sample Device Level VPN</string>
<key>PayloadOrganization</key>
<string>Palo Alto Networks</string>
<key>PayloadDescription</key>
<string>Profile Description</string>
<key>PayloadIdentifier</key>
<string>Sample Device Level VPN</string>
<key>PayloadType</key>
<string>Configuration</string>
<key>PayloadVersion</key>
<integer>1</integer>
<key>PayloadUUID</key>
<string>5436fc94-205f-7c59-0000-011c</string>
<key>PayloadRemovalDisallowed</key>
<false/>
</dict>
</plist>
```

Example: GlobalProtect iOS App App-Level VPN Configuration

The following example shows the XML configuration containing a VPN payload that you can use to verify the app-level VPN configuration of the GlobalProtect app for iOS.

```
<?xml version="1.0"
encoding="UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN" "http://www.apple.com/
DTDs/PropertyList-1.0.dtd">
<plist version="1.0">
<dict>
<key>PayloadContent</key>
<array>
<dict>
<key>PayloadDescription</key>
<string>Configures VPN settings, including authentication.</string>
<key>PayloadDisplayName</key>
<string>VPN (Sample App Level VPN)</string>
<key>PayloadIdentifier</key>
```


© 2017 Palo Alto Networks, Inc.

```
<string>Sample App Level VPN.vpn</string>
<key>PayloadOrganization</key>
<string>Palo Alto Networks</string>
<key>PayloadType</key>
<string>com.apple.vpn.managed.applayer</string>
<key>PayloadVersion</key>
<integer>1</integer>
<key>VPNUUID</key>
<string>cGFuU2FtcGxlIEFwcCBMZXZlbCBWUE52cG5TYW1wbGUgQXBwIExldmVsIFZQTg==</
string>
<key>SafariDomains</key>
<array>
<string>*.paloaltonetworks.com</string>
</array>
<key>PayloadUUID</key>
<string>54370008-205f-7c59-0000-01a1</string>
<key>UserDefinedName</key>
<string>Sample App Level VPN</string>
<key>Proxies</key>
<dict/>
<key>VPNType</key>
<string>VPN</string>
<key>VPNSubType</key>
<string>com.paloaltonetworks.GlobalProtect.vpnplugin</string>
<key>IPv4</key>
<dict>
<key>OverridePrimary</key>
<integer>0</integer>
</dict>
<key>VPN</key>
<dict>
<key>RemoteAddress</key>
<string>cademogp.paloaltonetworks.com</string>
<key>AuthName</key>
<string></string>
<key>OnDemandMatchAppEnabled</key>
<integer>1</integer>
<key>OnDemandEnabled</key>
<integer>1</integer>
<key>DisconnectOnIdle</key>
<integer>0</integer>
<key>AuthenticationMethod</key>
<string>Password</string>
</dict>
<key>VendorConfig</key>
<dict>
<key>OnlyAppLevel</key>
<integer>1</integer>
<key>AllowPortalProfile</key>
<integer>0</integer>
<key>FromAspen</key>
<integer>1</integer>
</dict>
</dict>
</array>
<key>PayloadDisplayName</key>
<string>Sample App Level VPN</string>
<key>PayloadOrganization</key>
<string>Palo Alto Networks</string>
<key>PayloadDescription</key>
<string>Profile Description</string>
<key>PayloadIdentifier</key>
```


© 2017 Palo Alto Networks, Inc.

```
<string>Sample App Level VPN</string>
<key>PayloadType</key>
<string>Configuration</string>
<key>PayloadVersion</key>
<integer>1</integer>
<key>PayloadUUID</key>
<string>5436fc94-205f-7c59-0000-011c</string>
<key>PayloadRemovalDisallowed</key>
<false/>
</dict>
</plist>
```

Configure the GlobalProtect App for Android

You can deploy and configure the GlobalProtect app on Android For Work devices from any third-party mobile device management (MDM) system supporting Android For Work App data restrictions.

On Android devices, traffic is routed through the VPN tunnel according to the access routes configured on the GlobalProtect gateway. From your third-party MDM that manages Android for Work devices, you can further refine the traffic that is routed though the VPN tunnel.

In an environment where the device is corporately owned, the device owner manages the entire device including all the apps installed on that device. By default, all installed apps can send traffic through the VPN tunnel according to the access routes defined on the gateway.

In a bring-your-own-device (BYOD) environment, the device is not corporately owned and uses a Work Profile to separate business and personal apps. By default only managed apps in the Work Profile can send traffic through the VPN tunnel according to the access routes defined on the gateway. Apps installed on the personal side of the device can not send traffic through the VPN tunnel set by the managed GlobalProtect app installed in the Work Profile.

To route traffic from an even smaller set of apps, you can enable Per-App VPN so that GlobalProtect only routes traffic from specific managed apps. For Per-App VPN, you can whitelist or blacklist specific managed apps from having their traffic routed through the VPN tunnel.

As part of the VPN configuration, you can also specify how the user connects to the VPN. When you configure the VPN connection method as **user-logon**, the GlobalProtect app will establish a connection automatically. When you configure the VPN connection method as **on-demand**, users can initiate a connection manually when attempting to connect to the VPN remotely.

The VPN connect method defined in the MDM takes precedence over the connect method defined in the GlobalProtect portal configuration.

Removing the VPN configuration automatically restores the GlobalProtect app to the original configuration settings.

To configure the GlobalProtect app for Android, configure the following Android App Restrictions.

Key	Value Type	Example
portal	String	**10.1.8.190**
username	String	**john**
password	String	**Passwd!234**

© 2017 Palo Alto Networks, Inc.

Key	Value Type	Example
certificate	String (in Base64)	**DAFDSaweEWQ23wDSAFD….**
client_certificate_passphrase	String	**PA$$W0RD$123**
app_list*	String	**whiltelist | blacklist: com.google.calendar; com.android.email; com.android.chrome**
connect_method	String	**user-logon | on-demand**
remove_vpn_config_via_ restriction	Boolean	**true | false**

*The `app_list` key specifies the configuration for Per-App VPN. Begin the string with either the whitelist or blacklist, and follow it with an array of app names separated by semicolon. The whitelist specifies the apps that will use the VPN tunnel for network communication. The network traffic for any other app that is not in the whitelist or expressly listed in the blacklist will not go through the VPN tunnel.

Example: Set VPN Configuration

```
private static String RESTRICTION_PORTAL
= "portal";
private static String RESTRICTION_USERNAME = "username";
private static String RESTRICTION_PASSWORD = "password";
private static String RESTRICTION_CONNECT_METHOD = "connect_method";
private static String RESTRICTION_CLIENT_CERTIFICATE
= "client_certificate";
private static String RESTRICTION_CLIENT_CERTIFICATE_PASSPHRASE
= "client_certificate_passphrase";
private static String RESTRICTION_APP_LIST = "app_list";
private static String RESTRICTION_REMOVE_CONFIG =
 "remove_vpn_config_via_restriction";

Bundle config = new Bundle();
config.putString(RESTRICTION_PORTAL, "192.168.1.1");
config.putString(RESTRICTION_USERNAME, "john");
config.putString(RESTRICTION_PASSWORD, "Passwd!234");
config.putString(RESTRICTION_CONNECT_METHOD, "user-logon");
config.putString(RESTRICTION_CLIENT_CERTIFICATE, "DAFDSaweEWQ23wDSAFD….");
config.putString(RESTRICTION_CLIENT_CERTIFICATE_PASSPHRASE,
"PA$$W0RD$123");
config.putString(RESTRICTION_APP_LIST,
 "whitelist:com.android.chrome;com.android.calendar");

DevicePolicyManager dpm = (DevicePolicyManager)
getSystemService(Context.DEVICE_POLICY_SERVICE);
dpm.setApplicationRestrictions(EnforcerDeviceAdminReceiver.getComponentName(this),
"com.paloaltonetworks.globalprotect", config);
```

Example: Remove VPN Configuration

```
Bundle config = new Bundle();
```


© 2017 Palo Alto Networks, Inc.

```
config.putBoolean(RESTRICTION_REMOVE_CONFIG, true );
DevicePolicyManager dpm = (DevicePolicyManager)
getSystemService(Context.DEVICE_POLICY_SERVICE);
dpm.setApplicationRestrictions(EnforcerDeviceAdminReceiver.
getComponentName(this),"com.paloaltonetworks.globalprotect",
config);
```

Host Information

Although you may have stringent security at your corporate network border, your network is really only as secure as the end devices that are accessing it. With today's workforce becoming more and more mobile, often requiring access to corporate resources from a variety of locations—airports, coffee shops, hotels—and from a variety of devices—both company-provisioned and personal—you must logically extend your network's security out to your endpoints to ensure comprehensive and consistent security enforcement. The GlobalProtect™ Host Information Profile (HIP) feature enables you to collect information about the security status of your end hosts—such as whether they have the latest security patches and antivirus definitions installed, whether they have disk encryption enabled, whether the device is jailbroken or rooted (mobile devices only), or whether it is running specific software you require within your organization, including custom applications—and base the decision as to whether to allow or deny access to a specific host based on adherence to the host policies you define.

The following topics provide information about the use of host information in policy enforcement. It includes the following sections:

> About Host Information on page 173
> Configure HIP-Based Policy Enforcement on page 177
> Collect Application and Process Data From Clients on page 184
> Block Device Access on page 190
> Configure Windows User-ID Agent to Collect Host Information on page 192

© 2017 Palo Alto Networks, Inc.

About Host Information

One of the jobs of the GlobalProtect agent is to collect information about the host it is running on. The agent then submits this host information to the GlobalProtect gateway upon successfully connecting. The gateway matches this raw host information submitted by the agent against any HIP objects and HIP profiles you have 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.

Using host information profiles for policy enforcement enables granular security that ensures that the remote hosts accessing your critical resources are adequately maintained and in adherence with your security standards before they are allowed access to your network resources. For example, before allowing access to your 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 only allows access to the application if the client system has encryption enabled. In addition, for clients 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 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).

- What Data Does the GlobalProtect Agent Collect? on page 173
- How Does the Gateway Use the Host Information to Enforce Policy? on page 175
- How Do Users Know if Their Systems are Compliant? on page 176
- How Do I Get Visibility into the State of the End Clients? on page 176

What Data Does the GlobalProtect Agent Collect?

By default, the GlobalProtect agent collects vendor-specific data about the end user security packages that are running on the computer (as compiled by the OPSWAT global partnership program) and reports this data to the GlobalProtect gateway for use in policy enforcement.

Because security software must continually evolve to ensure end user protection, your GlobalProtect gateway licenses also enable you to get dynamic updates for the GlobalProtect data file with the latest patch and software versions available for each package.

While the agent collects a comprehensive amount of data about the host it is running on, you may have additional software that you require your end-users to run in order to connect to your network or to access certain resources. In this case, you can define custom checks that instruct the agent to collect specific registry information (on Windows clients), preference list (plist) information (on Mac OS clients), or to collect information about whether or not specific services are running on the host.

The agent collects data about the following categories of information by default, to help to identify the security state of the host:

Table 6: Table: Data Collection Categories

Category	Data Collected
General	Information about the host itself, including the hostname, logon domain, operating system, client version, and, for Windows systems, the domain to which the machine belongs. *For Windows clients' domain, the GlobalProtect agent collects the domain defined for*

© 2017 Palo Alto Networks, Inc.

Category	Data Collected
	`ComputerNameDnsDomain`, which is the DNS domain assigned to the local computer or the cluster associated with the local computer. This data is what is displayed for the Windows clients' Domain in the HIP Match log details (Monitor > HIP Match).
Patch Management	Information about any patch management software that is enabled and/or installed on the host and whether there are any missing patches.
Firewall	Information about any client firewalls that are installed and/or enabled on the host.
Antivirus	Information about any antivirus software that is enabled and/or installed on the host, whether or not real-time protection is enabled, the virus definition version, last scan time, the vendor and product name. GlobalProtect uses OPSWAT technology to detect and assess third-party security applications on the endpoint. By integrating with the OPSWAT OESIS framework, GlobalProtect enables you to assess the compliance state of the endpoint. For example, you can define HIP objects and HIP profiles that verify the presence of a specific version of Antivirus software from a specific vendor on the endpoint and also ensure that it has the latest virus definition files.
Anti-Spyware	Information about any anti-spyware software that is enabled and/or installed on the host, whether or not real-time protection is enabled, the virus definition version, last scan time, the vendor and product name.
Disk Backup	Information about whether disk backup software is installed, the last backup time, and the vendor and product name of the software.
Disk Encryption	Information about whether disk encryption software is installed, which drives and/or paths are configured for encryption, and the vendor and product name of the software.
Data Loss Prevention	Information about whether data loss prevention (DLP) software is installed and/or enabled for the prevention sensitive corporate information from leaving the corporate network or from being stored on a potentially insecure device. This information is only collected from Windows clients.
Mobile Devices	Information about the mobile device, including the device name, logon domain, operating system, app version, and the mobile device network information to which the device is connected. In addition, GlobalProtect collects whether the device is rooted or jailbroken.  *To collect mobile device attributes and utilize them in HIP enforcement policies, GlobalProtect requires*

© 2017 Palo Alto Networks, Inc.

Category	Data Collected
	an MDM server. GlobalProtect currently supports HIP integration with the AirWatch MDM server. For devices managed by AirWatch, host information collected by the GlobalProtect app can be supplemented with additional information collected from the AirWatch service. Refer to Configure Windows User-ID Agent to Collect Host Information for a list of attributes that can be retrieved from AirWatch.

You can exclude certain categories of information from being collected on certain hosts (to save CPU cycles and improve client response time). To do this, you create a client configuration on the portal excluding the categories you are not interested in. For example, if you do not plan to create policy based on whether or not client systems run disk backup software, you can exclude that category and the agent will not collect any information about disk backup.

You can also choose to exclude collecting information from personal devices in order to allow for user privacy. This can include excluding device location and a list of apps installed on the device that are not managed by a third-party mobile device manager.

How Does the Gateway Use the Host Information to Enforce Policy?

While the agent gets the information about what information to collect from the client configuration downloaded from the portal, you define which host attributes you are interested in monitoring and/or using for policy enforcement by creating HIP objects and HIP profiles on the gateway(s):

- **HIP Objects**—Provide the matching criteria to filter out the host information you are interested in using to enforce policy from the raw data reported by the agent. For example, while the raw host data may include information about several antivirus packages that are installed on the client you may only be interested in one particular application that you require within your organization. In this case, you would create a HIP object to match the specific application you are interested in enforcing.

  The best way to determine what HIP objects you need is to determine how you will use the host information you collect to enforce policy. Keep in mind that the HIP objects themselves are merely building blocks that allow you to create the HIP profiles that are used in your security policies. Therefore, you may want to keep your objects simple, matching on one thing, such as the presence of a particular type of required software, membership in a specific domain, or the presence of a specific client OS. By doing this, you will have the flexibility to create a very granular (and very powerful) HIP-augmented policy.
- **HIP Profiles**—A collection of HIP objects that are to be evaluated together, either for monitoring or for security policy enforcement. When you create your HIP profiles, you can combine the HIP objects you previously created (as well as other HIP profiles) using Boolean logic such that when a traffic flow is evaluated against the resulting HIP profile it will either match or not match. If there is a match, the corresponding policy rule will be enforced; if there is not a match, the flow will be evaluated against the next rule, as with any other policy matching criteria.

Unlike a traffic log—which only creates a log entry if there is a policy match—the HIP Match log generates an entry whenever the raw data submitted by an agent matches a HIP object and/or a HIP profile you have defined. This makes the HIP Match log a good resource for monitoring the state of the hosts on your network over time—before attaching your HIP profiles to security policies—in order to help you determine exactly what policies you believe need enforcement. See Configure HIP-Based Policy Enforcement on page 177 for details on how to create HIP objects and HIP profiles and use them as policy match criteria.

© 2017 Palo Alto Networks, Inc.

How Do Users Know if Their Systems are Compliant?

By default, end users are not given any information about policy decisions that were made as a result of enforcement of a HIP-enabled security rule. However, you can enable this functionality by defining HIP notification messages to display when a particular HIP profile is matched and/or not matched.

The decision as to when to display a message (that is, whether to display it when the user's configuration matches a HIP profile in the policy or when it doesn't match it), depends largely on your policy and what a HIP match (or non-match) means for the user. That is, does a match mean they are granted full access to your network resources? Or does it mean they have limited access due to a non-compliance issue?

For example, consider the following scenarios:

- You create a HIP profile that matches if the required corporate antivirus and anti-spyware software packages are *not* installed. In this case, you might want to create a HIP notification message for users who match the HIP profile telling them that they need to install the software (and, optionally, providing a link to the file share where they can access the installer for the corresponding software).
- You create a HIP profile that matches if those same applications *are* installed, you might want to create the message for users who do not match the profile, and direct them to the location of the install package.

See Configure HIP-Based Policy Enforcement on page 177 for details on how to create HIP objects and HIP profiles and use in defining HIP notification messages.

How Do I Get Visibility into the State of the End Clients?

Whenever an end host connects to GlobalProtect, the agent presents its HIP data to the gateway. The gateway then uses this data to determine which HIP objects and/or HIP profiles the host matches. For each match, it generates a HIP Match log entry. Unlike a traffic log—which only creates a log entry if there is a policy match—the HIP Match log generates an entry whenever the raw data submitted by an agent matches a HIP object and/or a HIP profile you have defined. This makes the HIP Match log a good resource for monitoring the state of the hosts on your network over time—before attaching your HIP profiles to security policies—in order to help you determine exactly what policies you believe need enforcement.

Because a HIP Match log is only generated when the host state matches a HIP object you have created, for full visibility in to host state you may need to create multiple HIP objects to log HIP matches for hosts that are in compliance with a particular state (for security policy enforcement purposes) as well as hosts that are non-compliant (for visibility). For example, suppose you want to prevent a host that does not have Antivirus software installed from connecting to the network. In this case you would create a HIP object that matches hosts that have a particular Antivirus software installed. By including this object in a HIP profile and attaching it to the security policy rule that allows access from your VPN zone, you can ensure that only hosts that are protected with antivirus software can connect.

However, in this case you would not be able to see in the HIP Match log which particular hosts are not in compliance with this requirement. If you wanted to also see a log for hosts that do not have Antivirus software installed so that you can follow up with the users, you can also create a HIP object that matches the condition where the Antivirus software is not installed. Because this object is only needed for logging purposes, you do not need to add it to a HIP profile or attach it to a security policy rule.

© 2017 Palo Alto Networks, Inc.

Configure HIP-Based Policy Enforcement

To enable the use of host information in policy enforcement you must complete the following steps. For more information on the HIP feature, see About Host Information.

STEP 1 | Verify proper licensing for HIP checks.

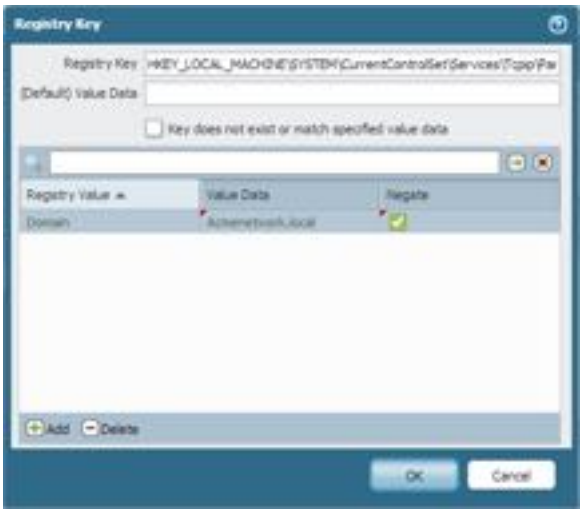


To use the HIP feature, you must have purchased and installed a GlobalProtect subscription license on each gateway that will perform HIP checks. To verify the status of your licenses on each portal and gateway, select **Device** > **Licenses**.

Contact your Palo Alto Networks Sales Engineer or Reseller if you do not have the required licenses. For more information on licensing, see About GlobalProtect Licenses.

STEP 2 | (**Optional**) Define any custom host information that you want the agent to collect. For example, if you have any required applications that are not included in the Vendor and/or Product lists for creating HIP objects, you could create a custom check that will allow you to determine whether that application is installed (has a corresponding registry or plist key) or is running (has a corresponding running process).

2 and 3 assume that you have already created a Portal Configuration. If you have not yet configured your portal, see Set Up Access to the GlobalProtect Portal for instructions.



1. On the firewall that is hosting your GlobalProtect portal, select **Network** > **GlobalProtect** > **Portals**.
2. Select your portal configuration to open the GlobalProtect Portal dialog.
3. Select the **Agent** tab and then select the agent configuration to which you want to add a custom HIP check, or click **Add** to create a new agent configuration.
4. Select the **Data Collection** tab.
5. Enable the option to **Collect HIP Data**.
6. Select **Custom Checks** and define the data you want to collect from hosts running this agent configuration as follows:

© 2017 Palo Alto Networks, Inc.

- **To collect information about specific registry keys**: On the **Windows** tab, **Add** the name of a **Registry Key** for which to collect data in the Registry Key area. Optionally, to restrict data collection to a specific Registry Value, **Add** and then define the specific Registry Value or values. Click **OK** to save the settings.
- **To collect information about running processes**: Select the appropriate tab (**Windows** or **Mac**) and then **Add** a process to the Process List. Enter the name of the process that you want the agent to collect information about.
- **To collect information about specific property lists**: On the **Mac** tab, click **Add** in the Plist section. Enter the **Plist** for which to collect data. Optionally, click **Add** to restrict the data collection to specific **Key** values. Click **OK** to save the settings.

7. If this is a new client configuration, complete the rest of the configuration as desired. For instructions, see Define the GlobalProtect Agent Configurations.
8. Click **OK** to save the client configuration.
9. **Commit** the changes.

STEP 3 | (**Optional**) Exclude categories from collection.

1. On the firewall that is hosting your GlobalProtect portal, select **Network** > **GlobalProtect** > **Portals**.
2. Select your portal configuration to open the GlobalProtect Portal dialog.
3. On the **Agent** tab, select the Agent configuration from which to exclude categories, or **Add** a new one.
4. Select **Data Collection**, and then verify that **Collect HIP Data** is enabled.
5. On the **Exclude Categories** tab, click **Add**. The Edit Exclude Category dialog displays.
6. Select the **Category** you want to exclude from the drop-down list.
7. (**Optional**) If you want to exclude specific vendors and/or products from collection within the selected category rather than excluding the entire category, click **Add**. You can then select the **Vendor** to exclude from the drop-down on the Edit Vendor dialog and, optionally, click **Add** to exclude specific products from that vendor. When you are done defining that vendor, click **OK**. You can add multiple vendors and products to the exclude list.
8. Repeat Step f and Step g for each category you want to exclude.
9. If this is a new client configuration, complete the rest of the configuration as desired. For more information on defining client configurations, see Define the GlobalProtect Agent Configurations.
10. Click **OK** to save the client configuration.
11. **Commit** the changes.

STEP 4 | Create the HIP objects to filter the raw host data collected by the agents.

The best way to determine what HIP objects you need is to determine how you will use the host information you collect to enforce policy. Keep in mind that the HIP objects themselves are merely building blocks that allow you to create the HIP profiles that are used in your security policies. Therefore, you may want to keep your objects simple, matching on one thing, such as the presence of a particular type of required software, membership in a specific domain, or the presence of a specific client OS. By doing this, you will have the flexibility to create a very granular (and very powerful) HIP-augmented policy.

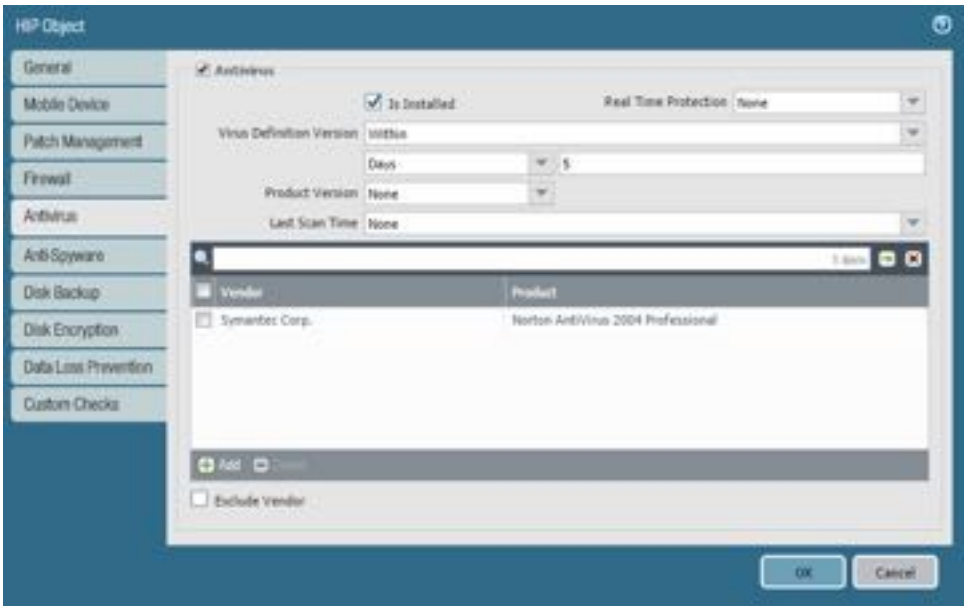


For details on a specific HIP category or field, refer to the online help.

1. On the gateway (or on Panorama if you plan to share the HIP objects among multiple gateways), select **Objects** > **GlobalProtect** > **HIP Objects** and click **Add**.
2. On the **General** tab, enter a **Name** for the object.
3. Select the tab that corresponds to the category of host information you are interested in matching against and select the check box to enable the object to match against the category. For example,

© 2017 Palo Alto Networks, Inc.

to create an object that looks for information about Antivirus software, select the **Antivirus** tab and then select the **Antivirus** check box to enable the corresponding fields. Complete the fields to define the desired matching criteria. For example, the following screenshot shows how to create an object that will match if the Symantec Norton AntiVirus 2004 Professional application is installed, has Real Time Protection enabled, and has virus definitions that have been updated within the last 5 days.



Repeat this step for each category you want to match against in this object. For more information, see Table: Data Collection Categories.

4. Click **OK** to save the HIP object.
5. Repeat these steps to create each additional HIP object you require.
6. **Commit** the changes.

STEP 5 | Create the HIP profiles that you plan to use in your policies.

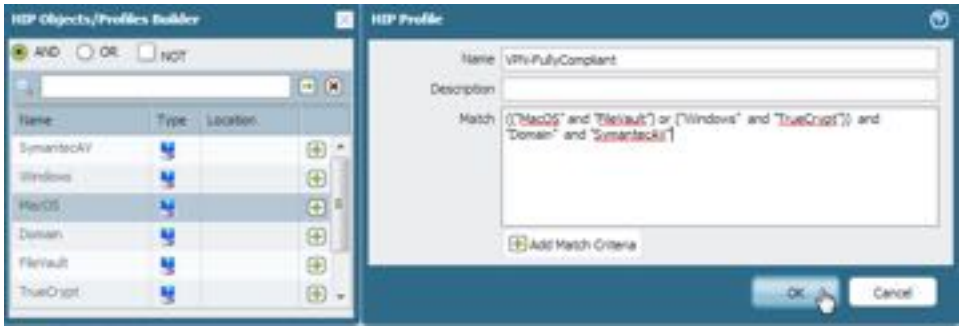
When you create your HIP profiles, you can combine the HIP objects you previously created (as well as other HIP profiles) using Boolean logic such that when a traffic flow is evaluated against the resulting HIP profile it will either match or not match. If there is a match, the corresponding policy rule will be enforced; if there is not a match, the flow will be evaluated against the next rule, as with any other policy matching criteria.

1. On the gateway (or on Panorama if you plan to share the HIP profiles among multiple gateways), select **Objects** > **GlobalProtect** > **HIP Profiles** and click **Add**.
2. Enter a descriptive **Name** for the profile and optionally a **Description**.
3. Click **Add Match Criteria** to open the HIP Objects/Profiles Builder.
4. Select the first HIP object or profile you want to use as match criteria and then click add ⊕ to move it over to the **Match** text box on the HIP Profile dialog. Keep in mind that if you want the HIP profile to evaluate the object as a match only when the criteria in the object is not true for a flow, select the **NOT** check box before adding the object.

© 2017 Palo Alto Networks, Inc.



5. Continue adding match criteria as appropriate for the profile you are building, making sure to select the appropriate Boolean operator radio button (**AND** or **OR**) between each addition (and, again, using the **NOT** check box when appropriate).
6. If you are creating a complex Boolean expression, you must manually add the parenthesis in the proper places in the **Match** text box to ensure that the HIP profile is evaluated using the logic you intend. For example, the following HIP profile will match traffic from a host that has either FileVault disk encryption (for Mac OS systems) or TrueCrypt disk encryption (for Windows systems) and also belongs to the required Domain, and has a Symantec antivirus client installed:



7. When you are done adding match criteria, click **OK** to save the profile.
8. Repeat these steps to create each additional HIP profile you require.
9. **Commit** the changes.

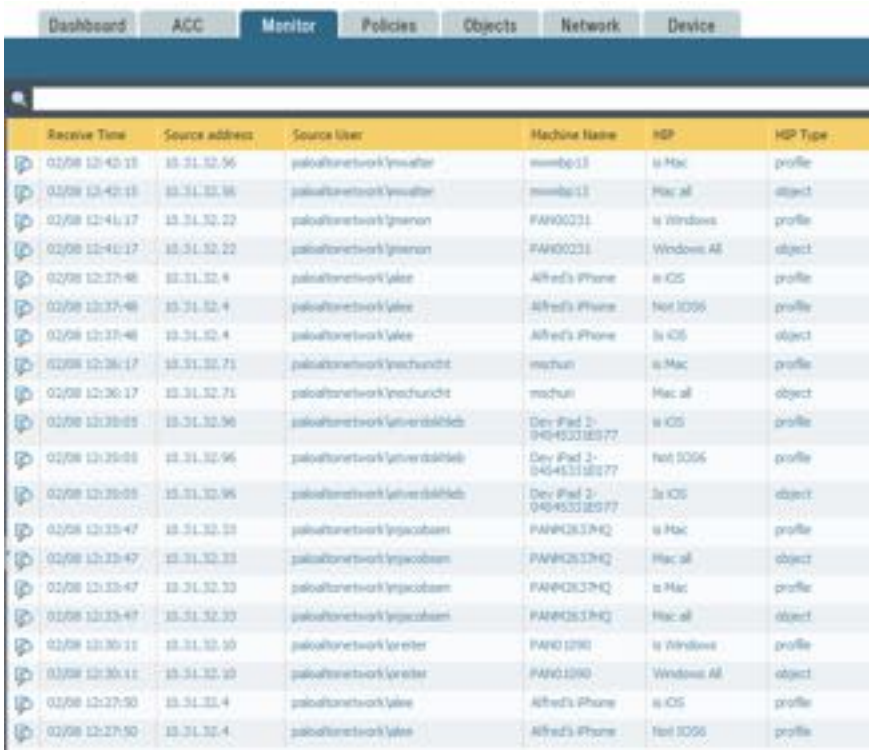
STEP 6 | Verify that the HIP objects and HIP profiles you created are matching your GlobalProtect client traffic as expected.



Consider monitoring HIP objects and profiles as a means to monitor the security state and activity of your host endpoints. By monitoring the host information over time you will be better able to understand where your security and compliance issues are and you can use this information to guide you in creating useful policy. For more details, see How Do I Get Visibility into the State of the End Clients?

On the gateway(s) that your GlobalProtect users are connecting to, select **Monitor** > **Logs** > **HIP Match**. This log shows all of the matches the gateway identified when evaluating the raw HIP data reported by the agents against the defined HIP objects and HIP profiles. Unlike other logs, a HIP match does not require a security policy match in order to be logged.

© 2017 Palo Alto Networks, Inc.



STEP 7 | Enable User-ID on the source zones that contain the GlobalProtect users that will be sending requests that require HIP-based access controls. You must enable User-ID even if you don't plan on using the user identification feature or the firewall will not generate any HIP Match logs entries.

1. Select **Network** > **Zones**.
2. Click on the **Name** of the zone in which you want to enable User-ID to open the Zone dialog.
3. Enable User ID by selecting the **Enabled** check box and then click **OK**.



STEP 8 | Create the HIP-enabled security rules on your gateway(s).

As a best practice, you should create your security rules and test that they match the expected flows based on the source and destination criteria as expected before adding your HIP profiles. By doing this you will also be better able to determine the proper placement of the HIP-enabled rules within the policy.

1. Select **Policies** > **Security** and select the rule to which you want to add a HIP profile.
2. On the **Source** tab, make sure the **Source Zone** is a zone for which you enabled User-ID in 7.
3. On the **User** tab, click **Add** in the **HIP Profiles** section and select the HIP profile(s) you want to add to the rule (you can add up to 63 HIP profiles to a rule).
4. Click **OK** to save the rule.
5. **Commit** the changes.



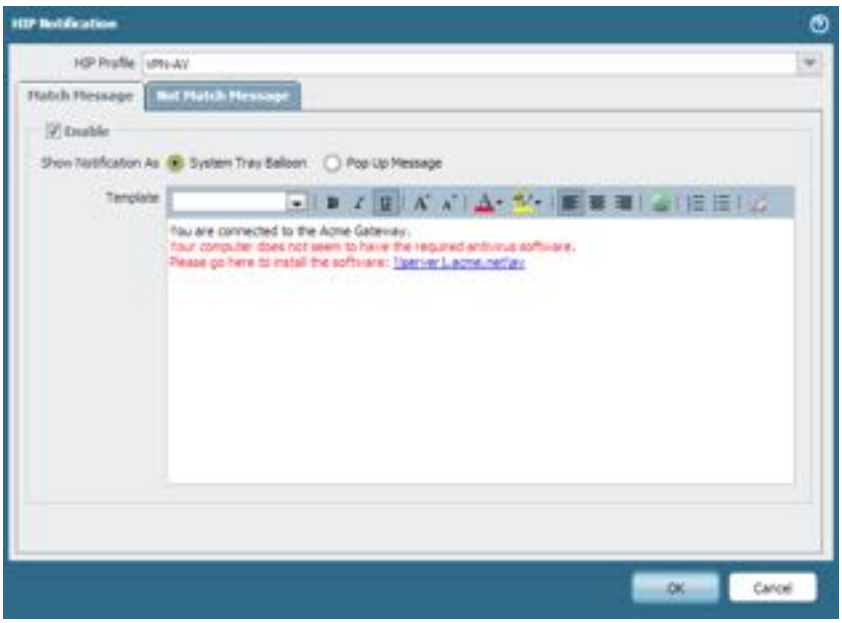
© 2017 Palo Alto Networks, Inc.

STEP 9 | Define the notification messages end users will see when a security rule with a HIP profile is enforced.

The decision as to when to display a message (that is, whether to display it when the user's configuration matches a HIP profile in the policy or when it doesn't match it), depends largely on your policy and what a HIP match (or non-match) means for the user. That is, does a match mean they are granted full access to your network resources? Or does it mean they have limited access due to a non-compliance issue?

For example, suppose you create a HIP profile that matches if the required corporate antivirus and anti-spyware software packages are not installed. In this case, you might want to create a HIP notification message for users who match the HIP profile telling them that they need to install the software. Alternatively, if your HIP profile matched if those same applications are installed, you might want to create the message for users who do not match the profile.

1. On the firewall that is hosting your GlobalProtect gateway(s), select **Network** > **GlobalProtect** > **Gateways**.
2. Select a previously-defined gateway configuration to open the GlobalProtect Gateway dialog.
3. Select **Client Configuration** > **HIP Notification** and then click **Add**.
4. Select the **HIP Profile** this message applies to from the drop-down.
5. Select **Match Message** or **Not Match Message**, depending on whether you want to display the message when the corresponding HIP profile is matched in policy or when it is not matched. In some cases you might want to create messages for both a match and a non-match, depending on what objects you are matching on and what your objectives are for the policy. For the Match Message, you can also enable the option to **Include matched application list in message** to indicate what applications triggered the HIP match.
6. Select the **Enable** check box and select whether you want to display the message as a **Pop Up Message** or as a **System Tray Balloon**.
7. Enter the text of your message in the Template text box and then click **OK**. The text box provides both a WYSIWYG view of the text and an HTML source view, which you can toggle between using the Source Edit icon. The toolbar also provides many options for formatting your text and for creating hyperlinks to external documents, for example to link users directly to the download URL for a required software program.



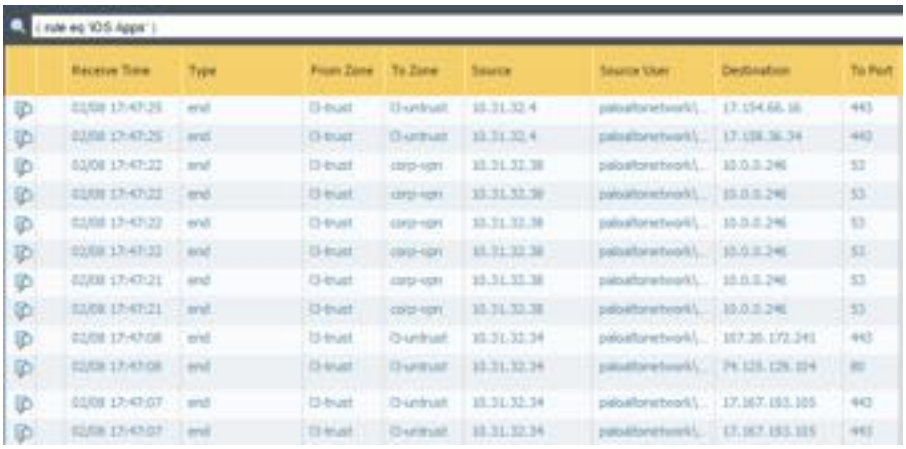
8. Repeat this procedure for each message you want to define.
9. **Commit** the changes.

© 2017 Palo Alto Networks, Inc.

STEP 10 | Verify that your HIP profiles are working as expected.

You can monitor what traffic is hitting your HIP-enabled policies using the Traffic log as follows:

1. From the gateway, select **Monitor** > **Logs** > **Traffic**.
2. Filter the log to display only traffic that matches the rule that has the HIP profile you are interested in monitoring attached. For example, to search for traffic that matches a security rule named “iOS Apps” you would enter **(rule eq 'iOS Apps')** in the filter text box as follows:



© 2017 Palo Alto Networks, Inc.

Collect Application and Process Data From Clients

The Windows Registry and Mac Plist can be used to configure and store settings and options for Windows and Mac operating systems, respectively. You can create a custom check that will allow you to determine whether an application is installed (has a corresponding registry or plist key) or is running (has a corresponding running process) on a Windows or Mac client. Enabling custom checks instructs the GlobalProtect agent to collect specific registry information (Registry Keys and Registry Key Values from Windows clients), preference list (plist) information (plist and plist keys from Mac OS clients). The data that you define to be collected in a custom check is included in the raw host information data collected by the GlobalProtect agent and then submitted to the GlobalProtect gateway when the agent connects.

To monitor the data collected with custom checks you can create a HIP object. You can then add the HIP object to a HIP profile to use the collected data to match to device traffic and enforce security rules. The gateway can use the HIP object (which matches to the data defined in the custom check) to filter the raw host information submitted by the agent. When the gateway matches the client data to a HIP object, a HIP Match log entry is generated for the data. A HIP profile allows the gateway to also match the collected data to a security rule. If the HIP profile is used as criteria for a security policy rule, the gateway will enforce that security rule on the matching traffic.

Use the following task to enable custom checks to collect data from Windows and Mac clients. This task includes the optional steps to create a HIP object and HIP profile for a custom check, if you would like to use client data as matching criteria for a security policy to monitor, identify, and act on traffic.



For more information on defining agent settings directly from the Windows registry or the global Mac plist, see Deploy Agent Settings Transparently.

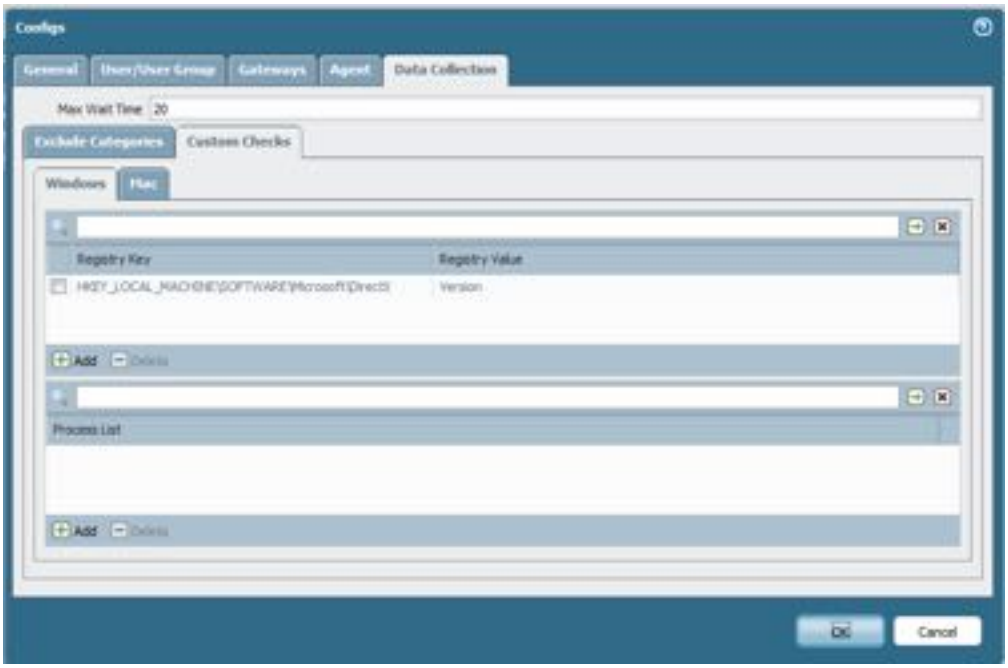
STEP 1 | Enable the GlobalProtect agent to collect Windows Registry information from Windows clients or Plist information from Mac clients. The type of information collected can include whether or not an application is installed on the client, or specific attributes or properties of that application.

This step enables the agent to report data on the applications and client settings. (5 and 6 will show you how to monitor and use the reported data to identify or take action on certain device traffic).

Collect data from a Windows client:

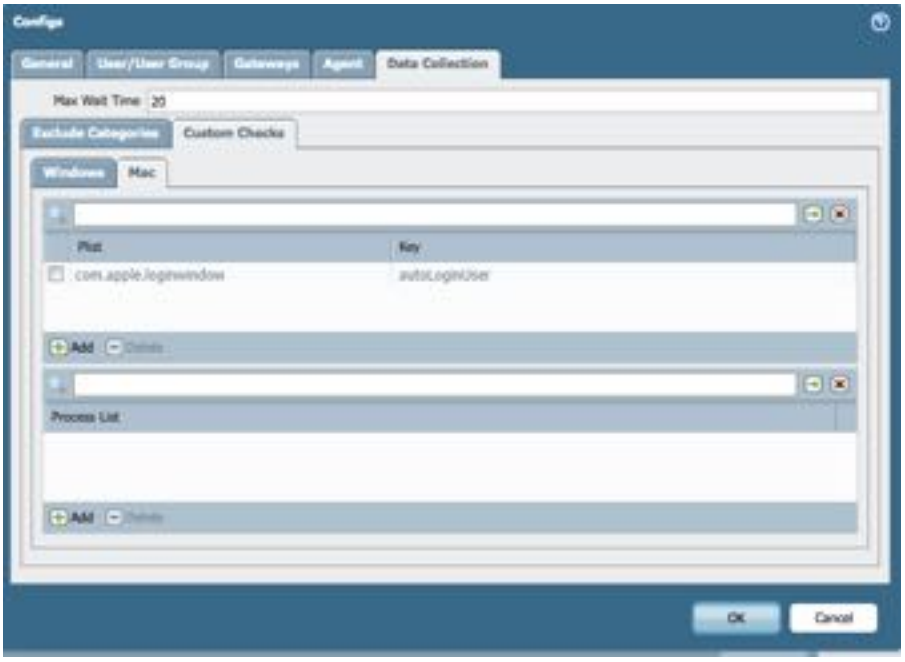
1. Select **Network** > **GlobalProtect** > **Portals** and then select the portal configuration you want to modify or **Add** a new one.
2. Select the **Agent** tab and then select the Agent configuration you want to modify or **Add** a new one.
3. Select **Data Collection**, and then verify that **Collect HIP Data** is enabled.
4. Select **Custom Checks** > **Windows**.
5. Add the Registry Key that you want to collect information about. If you want to restrict data collection to a value contained within that Registry Key, add the corresponding **Registry Value**.

© 2017 Palo Alto Networks, Inc.



Collect data from a Mac client:

1. Select **Network** > **GlobalProtect** > **Portals** and then select the portal configuration you want to modify or **Add** a new one.
2. Select the **Agent** tab and then select the Agent configuration you want to modify or **Add** a new one.
3. Select **Data Collection**, and then verify that **Collect HIP Data** is enabled.
4. Select **Custom Checks** > **Mac**.
5. Add the **Plist** that you want to collect information about and the corresponding Plist **Key** to determine if the application is installed:

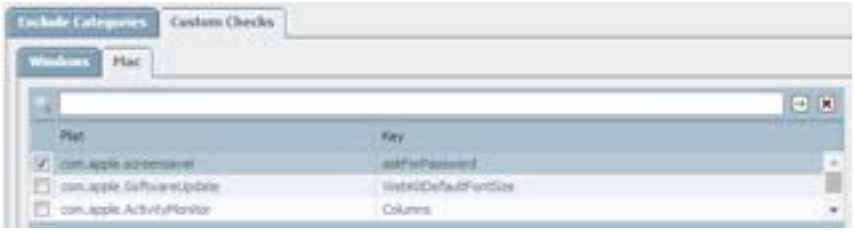


For example, **Add** the **Plist** `com.apple.screensaver` and the **Key** `askForPassword` to collect information on whether a password is required to wake the Mac client after the screen saver begins:

© 2017 Palo Alto Networks, Inc.



Confirm that the **Plist** and **Key** are added to the Mac custom checks:



STEP 2 | (**Optional**) Check if a specific process is running on the client.

1. Continue from 1 on the **Custom Checks** tab (**Network** > **GlobalProtect** > **Portals** > **<portal-config** > **Agent** > **<agent-config** > **Data Collection**) and select the **Windows** tab or **Mac** tab.
2. **Add** the name of the process that you want to collect information about to the **Process List**.

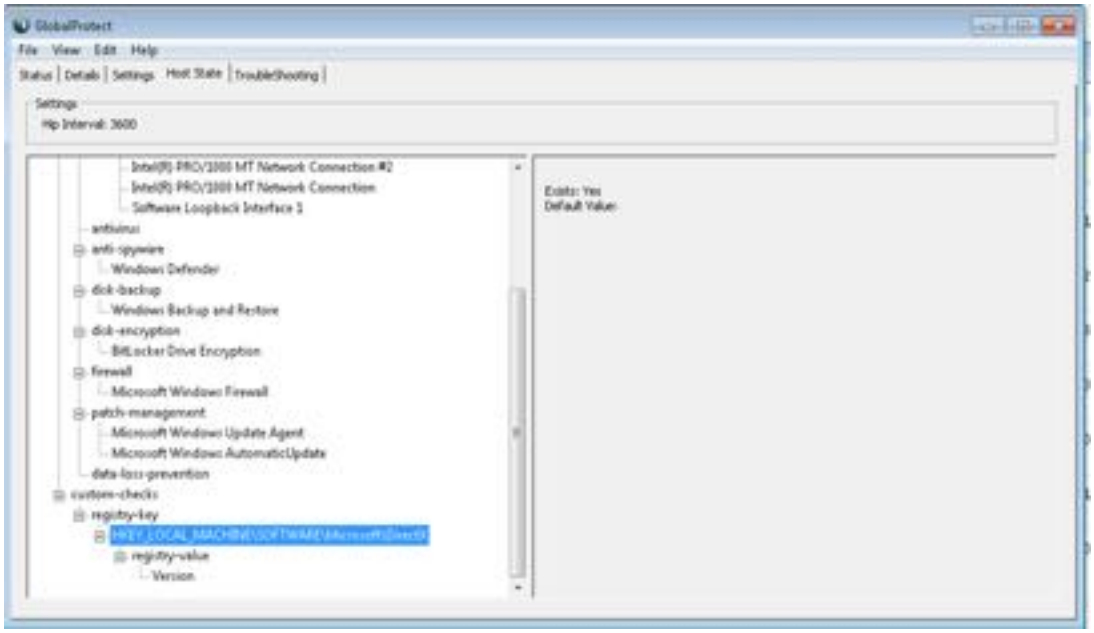
STEP 3 | Save the custom check.

Click **OK** and **Commit** the changes.

STEP 4 | Verify that the GlobalProtect agent is collecting the data defined in the custom check from the client.

For Windows clients:

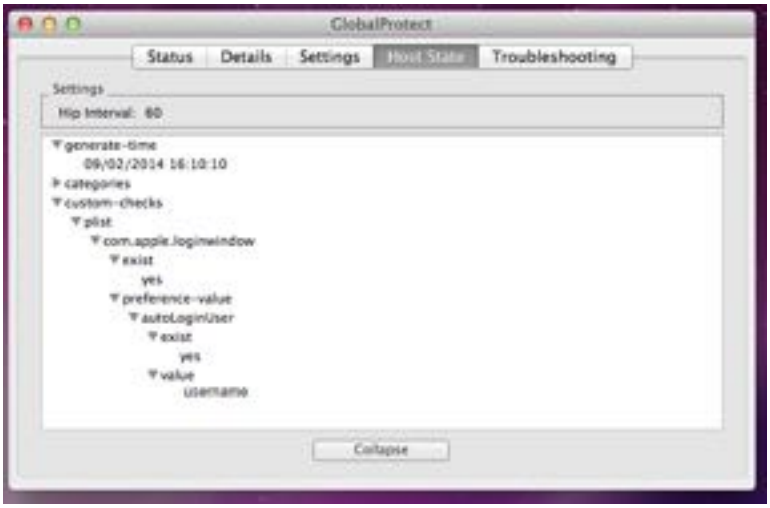
On the Windows client, double-click the GlobalProtect icon on the task bar and click the **Host State** tab to view the information that the GlobalProtect agent is collecting from the Mac client. Under the custom-checks dropdown, verify that the data that you defined for collection in 7 is displayed:



For Mac clients:

© 2017 Palo Alto Networks, Inc.

On the Mac client, click the GlobalProtect icon on the Menu bar, click **Advanced View**, and click **Host State** to view the information that the GlobalProtect agent is collecting for the Mac client. Under the custom-checks dropdown, verify that the data you defined for collection in 7 is displayed:



STEP 5 | (**Optional**) Create a HIP Object to match to a Registry Key (Windows) or Plist (Mac). This can allow you to filter the raw host information collected from the GlobalProtect agent in order to monitor the data for the custom check.

With a HIP object defined for the custom check data, the gateway will match the raw data submitted from the agent to the HIP object and a HIP Match log entry is generated for the data (**Monitor** > **HIP Match**).

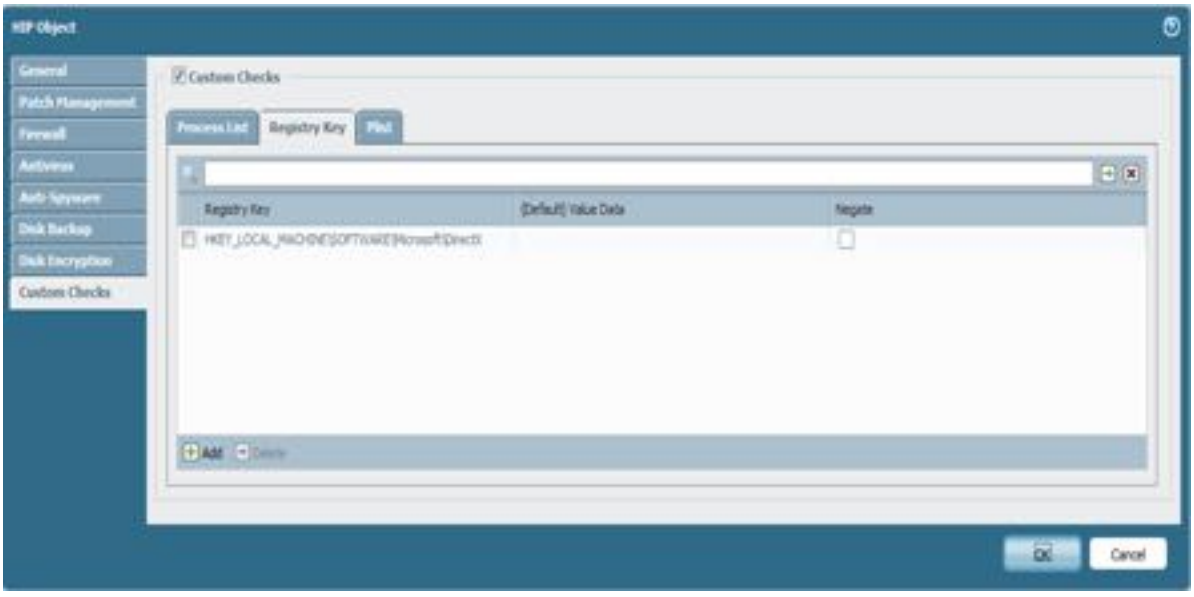
For Windows and Mac clients:

1. Select **Objects** > **GlobalProtect** > **HIP Objects** and **Add** a **HIP Object**.
2. Select and enable **Custom Checks**.

For Windows clients only:

1. To check Windows clients for a specific registry key, select **Registry Key** and **Add** the registry to match on. To only identify clients that do not have the specified registry key, select **Key does not exist or match the specified value data**.
2. To match on specific values within the Registry key, click **Add** and then enter the registry value and value data. To identify clients that explicitly do not have the specified value or value data, select the **Negate** check box.

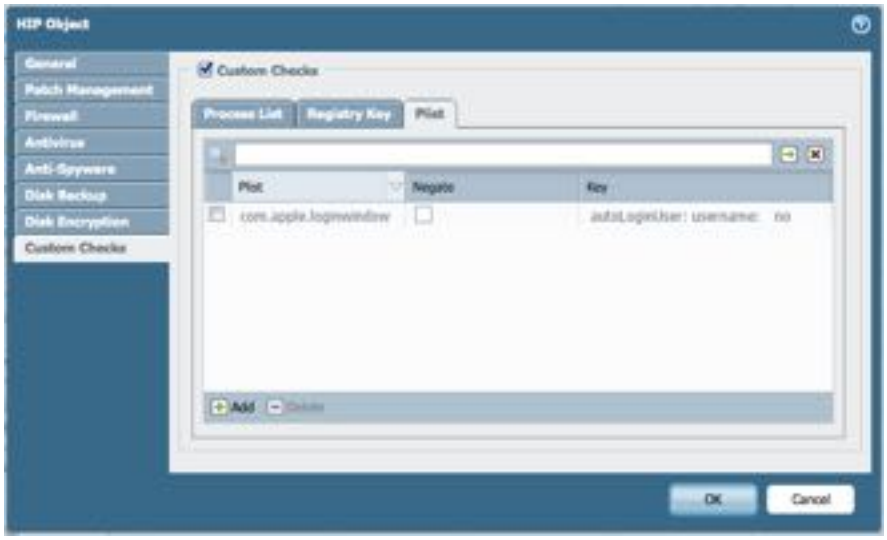
© 2017 Palo Alto Networks, Inc.



3. Click **OK** to save the HIP object. You can **Commit** to view the data in the **HIP Match** logs at the next device check-in or continue to 6.

For Mac clients only:

1. Select the **Plist** tab and **Add** and enter the name of the **Plist** for which you want to check Mac clients. (If instead, you want to match Mac clients that do not have the specified Plist, continue by selecting **Plist does not exist**).
2. (**Optional**) You can match traffic to a specific key-value pair within the Plist by entering the **Key** and the corresponding **Value** to match. (Alternatively, if you want to identify clients that do not have a specific Key and Value, you can continue by selecting **Negate** after adding populating the **Key** and **Value** fields).



3. Click **OK** to save the HIP object. You can **Commit** to view the data in the **HIP Match** logs at the next device check-in or continue to 6.

STEP 6 | (**Optional**) Create a HIP profile to allow the HIP object you created in 5 to be evaluated against traffic.

The HIP profile can be added to a security policy as an additional check for traffic matching that policy. When the traffic is matched to the HIP profile, the security policy rule will be enforced on the traffic.

© 2017 Palo Alto Networks, Inc.

For more details on creating a HIP profiles, see Configure HIP-Based Policy Enforcement.

1. Select **Objects** > **GlobalProtect** > **HIP Profile**.
2. Click **Add Match Criteria** to open the **HIP Objects/Profiles Builder**.
3. Select the **HIP object** you want to use as match criteria and then move it over to the **Match** box on the HIP Profile dialog.
4. When you have finished adding the objects to the new HIP profile, click **OK** and **Commit**.



STEP 7 | Add the HIP profile to a security policy so that the data collected with the custom check can be used to match to and act on traffic.

Select **Policies** > **Security**, and **Add** or modify a security policy. Go to the **User** tab to add a HIP profile to the policy. For more details on security policies components and using security policies to match to and act on traffic, see Security Policy.

© 2017 Palo Alto Networks, Inc.

Block Device Access

In the event that a user loses a device that provides GlobalProtect access to your network, that device is stolen, or a user leaves your organization, you can block the device from gaining access to the network by placing the device in a block list.

A block list is local to a logical network location (vsys, 1 for example) and can contain a maximum of 1,000 devices per location. Therefore, you can create separate device block lists for each location hosting a GlobalProtect deployments.

STEP 1 | Identify the host ID for the endpoints you want to block.

The host ID is a unique ID that GlobalProtect assigns to identify the host. The host ID value varies by device type:

- Windows—Machine GUID stored in the Windows registry (HKEY_Local_Machine\Software \Microsoft\Cryptography\MachineGuid)
- macOS—MAC address of the first built-in physical network interface
- Android—Android ID
- iOS—UDID
- Chrome—GlobalProtect assigned unique alphanumeric string with length of 32 characters

If you do not know the host ID, you can correlate the user-ID to the host ID in the HIP Match logs:

1. Select **Monitor** > **Logs** > **HIP Match**.
2. Filter the HIP match logs for the source user associated with the device.
3. Open the HIP match log and identify the host ID under **OS** > **Host ID** and optionally the hostname under **Host Information** > **Machine Name**.



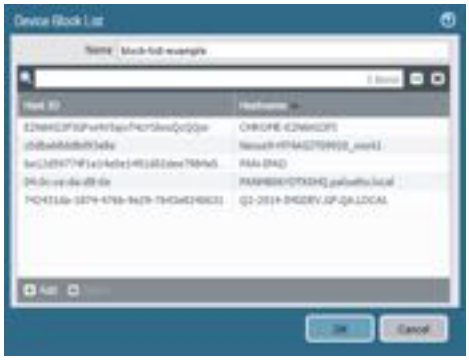
STEP 2 | Create a device block list.

© 2017 Palo Alto Networks, Inc.

You cannot use Panorama templates to push a device block list to firewalls.

1. Select **Network** > **GlobalProtect** > **Device Block List** and **Add** a device block list.
2. Enter a descriptive **Name** for the list.
3. For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available.

STEP 3 | Add a device to a block list.



1. **Add** devices. Enter the host ID (**required**) and hostname (**optional**) for a device you need to block.
2. **Add** additional devices, if needed.
3. Click **OK** to save and activate the block list.

The device list does not require a commit and is immediately active.

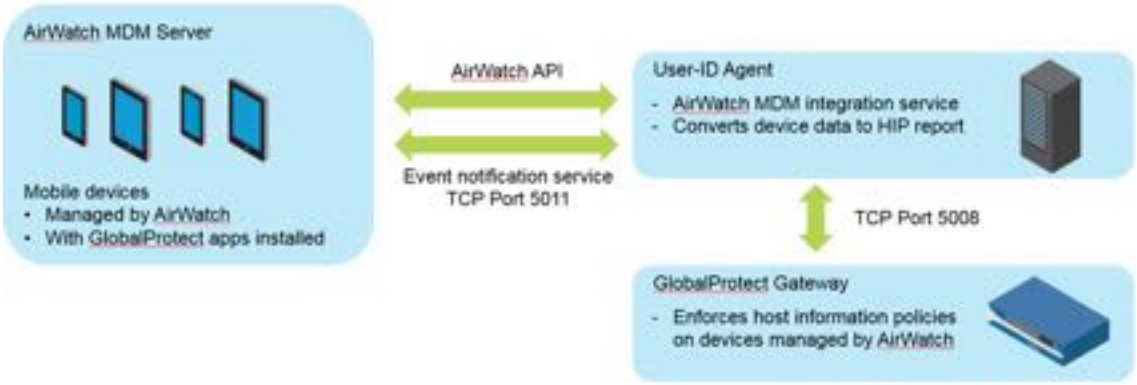
© 2017 Palo Alto Networks, Inc.

Configure Windows User-ID Agent to Collect Host Information

The Windows-based User-ID agent has been extended to support a new AirWatch MDM integration service. This service enables GlobalProtect to use the host information collected by the service to enforce HIP-based policies on devices managed by AirWatch. Running as part of the Windows-based User-ID agent, the AirWatch MDM integration service uses the AirWatch API to collect information from mobile devices that are managed by VMware AirWatch and translate this data into host information.

For Android devices managed by AirWatch, this feature supports Android for Work devices, but it does not support other types of Android devices.

- MDM Integration Overview on page 192
- Information Collected on page 192
- System Requirements on page 194
- Configure GlobalProtect to Retrieve Host Information on page 194
- Troubleshoot the MDM Integration Service on page 197



MDM Integration Overview

MDM integration service included with the Windows-based User-ID agent does a full HIP query to the AirWatch MDM server to get the complete host information for a device. GlobalProtect apps on the devices also send HIP information to the gateway and the gateway merges HIP information from the GlobalProtect apps and the MDM integration service. When a mobile device running the GlobalProtect app is connected to a GlobalProtect gateway, GlobalProtect can apply security policies with host information profiles.

You configure the MDM integration service to fetch AirWatch device information at regular intervals and push this information to GlobalProtect gateways. In addition, the service can monitor AirWatch event notifications and fetch updated device information when AirWatch events occur (for example, device enrollment, device wipe, and compliance changes).

Information Collected

The following table shows how information collected from devices managed by AirWatch are translated into HIP report attributes. The mapping is done automatically.

© 2017 Palo Alto Networks, Inc.

AirWatch Attributes	HIP Report Attributes
Device Information	
SerialNumber	serial-number
MacAddress	wifimac
Imei	IMEI
OperatingSystem	version
Model	model
DeviceFriendlyName	devname
IsSupervised	supervised
Udid (Unique Device Identifier)	udid
UserName	user
LastEnrolledOn	enroll-time
Platform	os
EnrollmentStatus	managed-by-mdm
LastSeen	last-checkin-time
ComplianceStatus (User-ID agent 8.0.3 and later)	Compliant NonCompliant NotAvailable
Ownership (User-ID agent 8.0.3 and later)	E—Employee Owned C—Corporate-Dedicated S—Corporate-Shared
Security Information	
DataProtectionEnabled	disk-encrypted
IsPasscodePresent	passcode-set
IsPasscodeCompliant	passcode-compliant
Network Information	
DataRoamingEnabled	data-roaming
GPS Coordinates	

© 2017 Palo Alto Networks, Inc.

AirWatch Attributes	HIP Report Attributes
Latitude	latitude
Longitude	longitude
SampleTime	last-location-time
Application Details	
ApplicationName	appname
Version	version
ApplicationIndentifier	package

System Requirements

AirWatch MDM integration service requires the following software:

Software	Minimum Supported Version
User-ID Agent	8.0.1
PAN-OS	7.0
GlobalProtect App for Android	4.0
GlobalProtect App for iOS	4.0.1
AirWatch Server	8.4.7.0
Windows Server	2008 and 2012 2016 with User-ID Agent 8.0.4 and PAN-OS 8.0.4 8.0.4

Configure GlobalProtect to Retrieve Host Information

Follow these instructions to configure GlobalProtect to retrieve host information from devices managed by AirWatch.

STEP 1 | Install and configure the Windows-based User-ID agent. The User-ID agent must be in a location that enables secure connections to the VMware AirWatch Mobile Device Management (MDM) system.

For more information, see Install the User-ID Agent. The AirWatch MDM integration service is included with the PAN-OS Windows-based User-ID agent.

STEP 2 | Configure SSL authentication between the Windows-based User-ID agent and the GlobalProtect gateway.

© 2017 Palo Alto Networks, Inc.

When you configure SSL authentication, make sure:

- The server certificate configured on the Windows-based User-ID Agent has the same Common Name (CN) as the hostname/IP address of the User-ID Agent host.
- The server certificate is trusted by the firewall (included in the trusted CA list in the MDM configuration on the firewall).
- The root certificate authority (CA) certificate of the MDM client certificate configured on the firewall must be imported into Windows trust store of the Windows server.

1. Obtain a server certificate and private key for authentication between the Windows-based User-ID agent and the GlobalProtect gateway. The certificate bundle must be in PEM format that contains a PEM certificate, full certificate chain, and private key.
2. Open the Windows-based User-ID agent and select **Server Certificate**.
3. **Add** the server certificate.

- **Browse** to the certificate file and **Open** the file to upload the certificate to the Windows-based User-ID agent.
- Enter a **Private Key Password** for the certificate.
- Click **OK.**

The agent verifies the certificate is valid and stores the encryption password of the private key in the host machine's Windows credential store.

If installation is successful, detailed information about the certificate (including common name, expiration date, and issuer) appears on the **Server Certificate** tab.

1. Restart the Windows-based User-ID agent.

STEP 3 | Configure the MDM integration service on the Windows-based User-ID agent.

1. Select **MDM Integration** in the Windows-based User-ID agent.
2. In **Gateway Connection TCP Port**, specify a port for TCP communications. The Windows-based User-ID agent listens at this port for all MDM-related messages. The default port is 5008. To change the port, specify a number from 1 to 65535.
3. On the **Setup** tab, click **Edit**.
4. Choose **AirWatch** for the **MDM Vendor**.

STEP 4 | Specify the **MDM Event Notification** settings to monitor and collect AirWatch events (for example, device enrollment, device wipe, and compliance changes). When an event occurs, the MDM integration service fetches the updated device information from the AirWatch API and pushes this information to all configured GlobalProtect gateways.



For MDM Event Notification, make sure the values you enter here are also configured in the AirWatch console under Groups & Settings > All Settings > System > Advanced > API > Event Notifications.

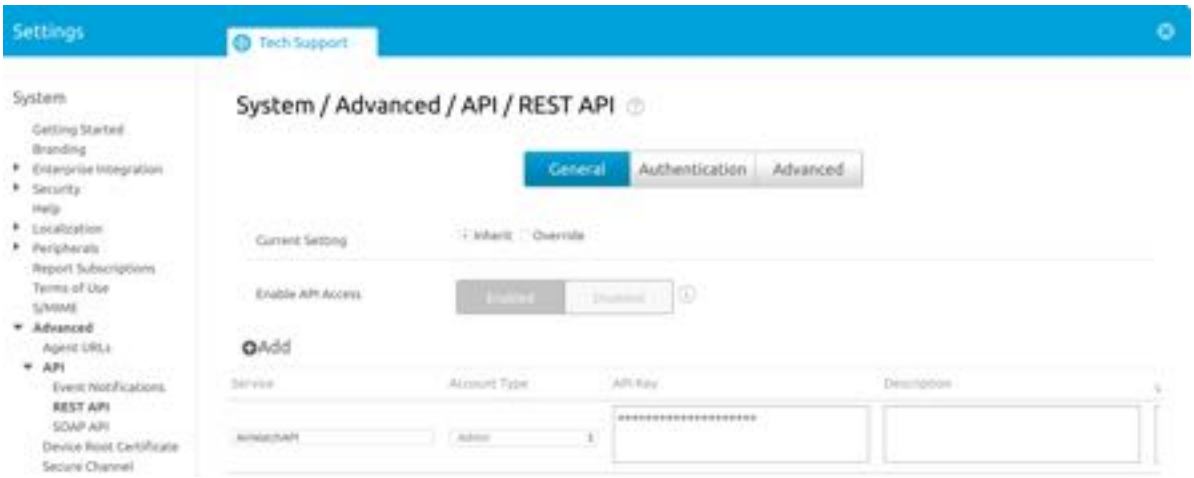
© 2017 Palo Alto Networks, Inc.



- Set the **TCP Port** for communicating with the event notification service. Use this format: **`http://<external_hostname>/<ip_address>:<port>`** where **`<ip-address>`** is the IP address for the MDM integration service. The default port is 5011. To change the port, specify a number from 1 to 65535.
- For event notification, enter the **Username** and **Password** credentials needed to authenticate incoming requests.
- Enter the **Permitted IP** addresses to access MDM events. This is a comma-separated list of IP addresses from where MDM events are posted. For example, the IP address of the AirWatch server. Contact your AirWatch Support team for guidance on which IP addresses to specify.

STEP 5 | Add **MDM API Authentication** settings to connect with the AirWatch API.

- Enter the **Server Address** of the AirWatch MDM server to which the Windows-based User-ID agent will connect. For example, **`api.awmdm.com`**.
- Enter the **Username** and **Password** credentials needed to access the AirWatch MDM API.
- Enter the **Tenant Code.**This is a unique hexadecimal code number required to access the AirWatch MDM API. On the AirWatch console, you can find the tenant code at **System** > **Advanced** > **API** > **REST API** > **API Key**.



- Enter the **Mobile Device State Retrieval Interval**. This setting controls how often host information is retrieved from devices managed by AirWatch. The default is 30 minutes. To change the interval, specify a number from 1 to 600.

STEP 6 | **Commit** your changes.

© 2017 Palo Alto Networks, Inc.

STEP 7 | Click **Test Connection** to make sure the Windows-based User-ID agent can connect to the AirWatch API.

STEP 8 | Configure the GlobalProtect gateway to communicate with the MDM integration service to retrieve the HIP reports for the devices managed by AirWatch.

1. In the PAN-OS web interface, select **Network** > **GlobalProtect** > **MDM**.
2. **Add** the following information about the MDM integration service.

- **Name**—Enter a name for the MDM integration service (up to 31 characters). The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
- (**Optional**) Select the virtual system to which the gateway belongs.
- **Server**—Enter the IP address or FQDN of the interface on the Airwatch MDM integration service where the gateway connects to retrieve HIP reports. Ensure that you have a service route to this interface.
- **Connection Port**—Enter the connection port where the MDM integration service listens for HIP report requests. The default port is 5008. To change the port, specify a number from 1 to 65535.
- **Client Certificate**—Choose the client certificate for the gateway to present to the MDM integration service when it establishes an HTTPS connection. You can choose a client certificate from the drop down, or import a new client certificate. The **Certificate Purpose** must indicate that it is a client authentication certificate.



The root certificate authority (CA) certificate of the client certificate must be imported into the Windows trust store of the Windows server where the User-ID Agent is installed.

1. **Add** the root CA certificate associated with the server certificate installed on the MDM integration service host. You need both the root CA certificate and the server certificate to establish a secure connection between the gateway and the MDM integration service. You can choose a root CA certificate from the drop down, or *Import* a new certificate.
2. Click **OK**.
3. **Commit** your changes.

STEP 9 | Check your connection to make sure AirWatch device data is transferred to GlobalProtect.

1. Open the Windows-based User-ID agent and select **MDM Integration** > **Mobile Devices**. You should see a list of unique device IDs and user names for all the devices managed by AirWatch.
2. (**Optional**) You can **Filter** the list to find a specific **Mobile Device**.
3. (**Optional**). Select a device in the list of device IDs and click **Retrieve Device State** to extract the latest information about the device and see how it maps to host information profiles on the GlobalProtect gateway.

Troubleshoot the MDM Integration Service

Follow these instructions if you have trouble with event notifications or trouble authenticating to the AirWatch REST API.

- Event notifications from the AirWatch MDM server are not received by the MDM integration service.
  1. Set the **Debug** option (in the **File** menu) to **Debug** or **Verbose**.
  2. Go the User-ID agent installation folder on the Windows server and open the `MaDebug` file. Look for messages similar to the following:

```
The address x.x.x.x
is not in the permitted ip list for event notifications.
```


© 2017 Palo Alto Networks, Inc.

3. Add this IP address as a **Permitted IP** address (**MDM Integration** > **Setup** > **Permitted IP**).

- Authentication to the Airwatch REST API is unsuccessful.

  Make sure that:

  - The credentials used for the MDM integration service to authenticate to the AirWatch MDM service are valid.
  - The user account used to access the Airwatch REST API has API access permissions and read-only permissions (at minimum) to data for the mobile devices and users managed by AirWatch.
  - The **Tenant Code** (API key) is correctly associated with the user account. Remove all unused API keys.

GlobalProtect Quick Configs

The following sections provide step-by-step instructions for configuring some common GlobalProtect™ deployments:

- Remote Access VPN (Authentication Profile)
- Remote Access VPN (Certificate Profile)
- Remote Access VPN with Two-Factor Authentication
- Always On VPN Configuration
- Remote Access VPN with Pre-Logon
- GlobalProtect Multiple Gateway Configuration
- GlobalProtect for Internal HIP Checking and User-Based Access
- Mixed Internal and External Gateway Configuration
- Live KB: Active Directory Password Changes

© 2017 Palo Alto Networks, Inc.

Remote Access VPN (Authentication Profile)

In the GlobalProtect VPN for Remote Access, the GlobalProtect portal and gateway are configured on ethernet1/2, so this is the physical interface where GlobalProtect clients connect. After a client connects and the portal and gateway authenticates it, the client establishes a VPN tunnel from its virtual adapter, which has been assigned an address in the IP address pool associated with the gateway tunnel.2 configuration—10.31.32.3-10.31.32.118 in this example. Because GlobalProtect VPN tunnels terminate in a separate corp-vpn zone, you have visibility into the VPN traffic as well as the ability to customize security policy for remote users.

Watch the video.

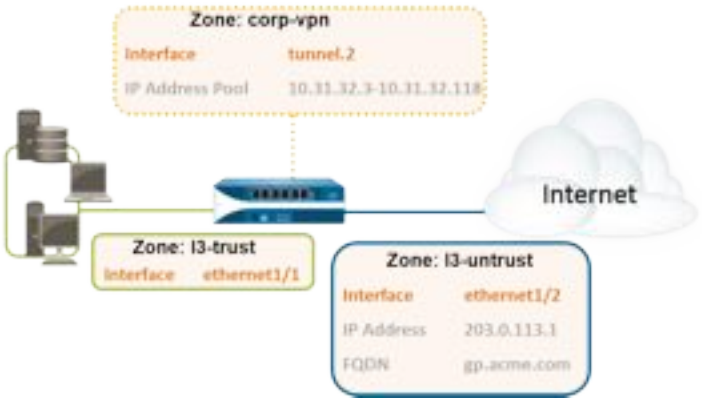


Figure 3: GlobalProtect VPN for Remote Access

The following procedure provides the configuration steps for this example. You can also watch the video.

STEP 1 | Create Interfaces and Zones for GlobalProtect.

Use the default virtual router for all interface configurations to avoid having to create inter-zone routing.

- Select **Network** > **Interfaces** > **Ethernet** and configure ethernet1/2 as a Layer 3 Ethernet interface with IP address 203.0.113.1 and assign it to the l3-untrust zone and the default virtual router.
- Create a DNS "A" record that maps IP address 203.0.113.1 to gp.acme.com.
- Select **Network** > **Interfaces** > **Tunnel** and add the tunnel.2 interface and add it to a new zone called corp-vpn. Assign it to the default virtual router.
- Enable User Identification on the corp-vpn zone.

STEP 2 | Create security policy to enable traffic flow between the corp-vpn zone and the l3-trust zone to enable access to your internal resources.

1. Select **Policies** > **Security** and then **Add** a new rule.
2. For this example, you would define the rule with the following settings:
   - Name—VPN Access
   - Source Zone—corp-vpn
   - Destination Zone—l3-trust

© 2017 Palo Alto Networks, Inc.



STEP 3 | Obtain a server certificate for the interface hosting the GlobalProtect portal and gateway using one of the following methods:

- (**Recommended**) Import a server certificate from a well-known, third-party CA.
- Use the root CA on the portal to generate a self-signed server certificate.

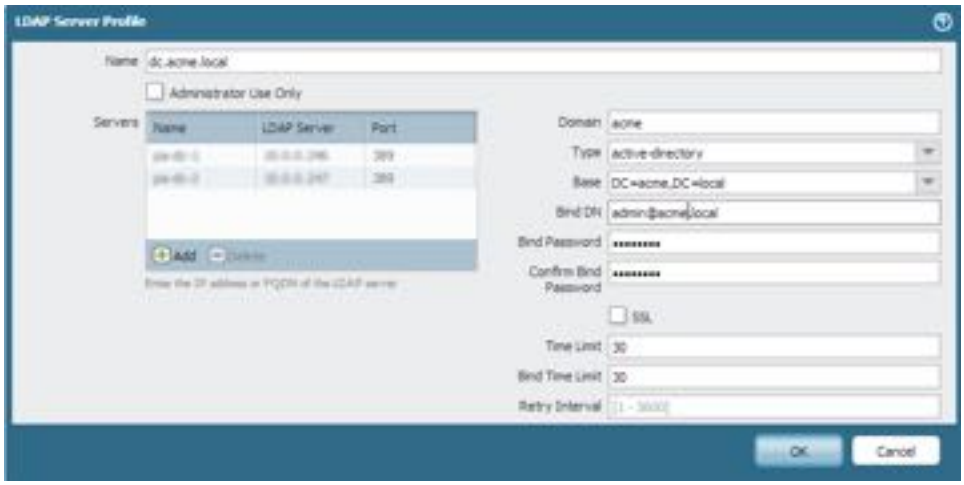
Select **Device** > **Certificate Management** > **Certificates** to manage certificates as follows:

- Obtain a server certificate. Because the portal and gateway are on the same interface, the same server certificate can be used for both components.
- The CN of the certificate must match the FQDN, `gp.acme.com`.
- To enable clients to connect to the portal without receiving certificate errors, use a server certificate from a public CA.

STEP 4 | Create a server profile.

The server profile instructs the firewall how to connect to the authentication service. Supported methods are Local, RADIUS, Kerberos, SAML, and LDAP authentication. This example shows an LDAP authentication profile for authenticating users against the Active Directory.

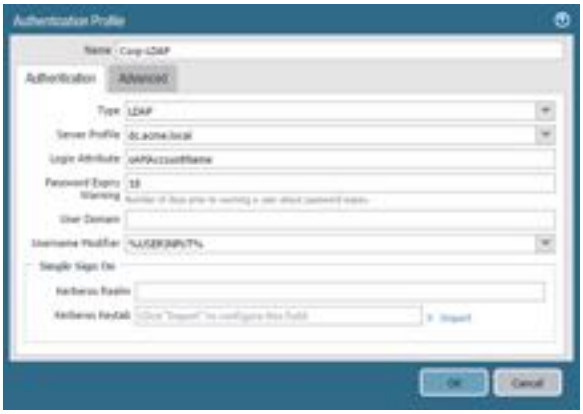
Create the server profile for connecting to the LDAP server (**Device** > **Server Profiles** > **LDAP**).



STEP 5 | (Optional) Create an authentication profile.

Attach the server profile to an authentication profile (**Device** > **Authentication Profile**).

© 2017 Palo Alto Networks, Inc.



STEP 6 | Configure a GlobalProtect Gateway.

Select **Network** > **GlobalProtect** > **Gateways** and add the following configuration:

Interface—**ethernet1/2**

IP Address—**203.0.113.1**

Server Certificate—**GP-server-cert.pem issued by GoDaddy**

Authentication Profile—**Corp-LDAP**

Tunnel Interface—**tunnel.2**

IP Pool—**10.31.32.3 - 10.31.32.118**

STEP 7 | Configure the GlobalProtect Portals.

Select **Network** > **GlobalProtect** > **Portals** and add the following configuration:

1. Set Up Access to the GlobalProtect Portal. This example uses the following settings:

   Interface—**ethernet1/2**

   IP Address—**203.0.113.1**

   Server Certificate—**GP-server-cert.pem issued by GoDaddy**

   Authentication Profile—**Corp-LDAP**
2. Define the GlobalProtect Client Authentication Configurations using the following settings:

   Connect Method—**On-demand** (Manual user initiated connection)

   External Gateway Address—**gp.acme.com**

STEP 8 | Deploy the GlobalProtect Agent Software.

Select **Device** > **GlobalProtect Client**.

In this example, use the procedure to Host Agent Updates on the Portal.

STEP 9 | (**Optional**) Enable use of the GlobalProtect mobile app.

Purchase and install a GlobalProtect subscription (**Device** > **Licenses**) to enable use of the app.

STEP 10 | Save the GlobalProtect configuration.

Click **Commit**.

© 2017 Palo Alto Networks, Inc.

Remote Access VPN (Certificate Profile)

With certificate authentication, the client must present a valid client certificate that identifies the user to the GlobalProtect portal or gateway. In addition to the certificate itself, the portal or gateway can use a certificate profile to determine whether the client that sent the certificate is the client to which the certificate was issued.

When a client certificate is the only means of authentication, the certificate that the client presents must contain the username in one of the certificate fields; typically the username corresponds to the common name (CN) in the Subject field of the certificate.

Upon successful authentication, the GlobalProtect agent establishes a VPN tunnel with the gateway and is assigned an IP address from the IP pool in the gateway's tunnel configuration. To support user-based policy enforcement on sessions from the corp-vpn zone, the username from the certificate is mapped to the IP address that the gateway assigned. Also, if a security policy requires a domain name in addition to user name, the specified domain value in the certificate profile is appended to the username.

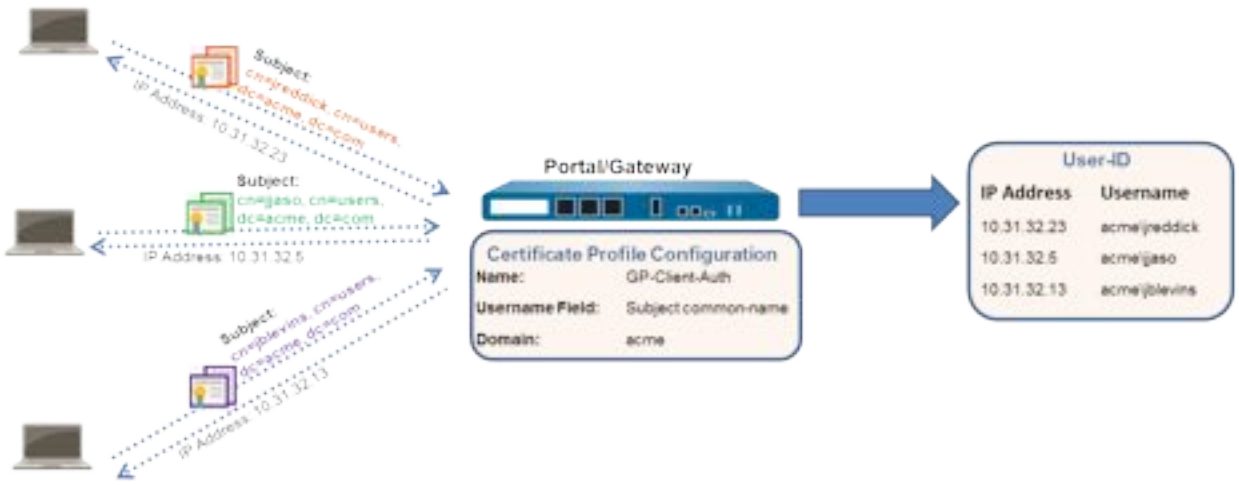


Figure 4: GlobalProtect Client Certificate Authentication Configuration

This quick configuration uses the same topology as GlobalProtect VPN for Remote Access. The only configuration difference is that instead of authenticating users against an external authentication server, this configuration uses client certificate authentication only.

STEP 1 | Create Interfaces and Zones for GlobalProtect.

Use the default virtual router for all interface configurations to avoid having to create inter-zone routing.

- Select **Network** > **Interfaces** > **Ethernet** and configure ethernet1/2 as a Layer 3 Ethernet interface with IP address **`203.0.113.1`** and assign it to the l3-untrust security zone and the default virtual router.
- Create a DNS "A" record that maps IP address **`203.0.113.1`** to **`gp.acme.com`**.
- Select **Network** > **Interfaces** > **Tunnel**.
- Add tunnel.2 interface to a new zone called `corp-vpn`. Assign the interface to the default virtual router.
- Enable User Identification on the corp-vpn zone.

© 2017 Palo Alto Networks, Inc.

STEP 2 | Create security policy to enable traffic flow between the corp-vpn zone and the l3-trust zone to enable access to your internal resources.

1. Select **Policies** > **Security** and then **Add** a new rule.
2. For this example, you would define the rule with the following settings:
   - Name—**VPN Access**
   - Source Zone—**corp-vpn**
   - Destination Zone—**l3-trust**



STEP 3 | Obtain a server certificate for the interface hosting the GlobalProtect portal and gateway using one of the following methods:

- (**Recommended**) Import a server certificate from a well-known, third-party CA.
- Use the root CA on the portal to generate a self-signed server certificate.

Select **Device** > **Certificate Management** > **Certificates** to manage certificates as follows:

- Obtain a server certificate. Because the portal and gateway are on the same interface, the same server certificate can be used for both components.
- The CN of the certificate must match the FQDN, `gp.acme.com`.
- To enable clients to connect to the portal without receiving certificate errors, use a server certificate from a public CA.

STEP 4 | Issue client certificates to GlobalProtect clients and endpoints.

1. Use your enterprise PKI or a public CA to issue a unique client certificate to each GlobalProtect user.
2. Install certificates in the personal certificate store on the endpoints.

STEP 5 | Create a client certificate profile.

1. Select **Device** > **Certificate Management** > **Certificate Profile**, click **Add** and enter a profile **Name** such as **GP-client-cert**.
2. Select **Subject** from the **Username Field** drop-down.
3. Click **Add** in the CA Certificates section, select the **CA Certificate** that issued the client certificates, and click **OK** twice.

STEP 6 | Configure a GlobalProtect Gateway.

See the topology diagram shown in GlobalProtect VPN for Remote Access.

Select **Network** > **GlobalProtect** > **Gateways** and add the following configuration:

Interface—**ethernet1/2**

IP Address—**203.0.113.1**

Server Certificate—**GP-server-cert.pem issued by GoDaddy**

Certificate Profile—**GP-client-cert**

Tunnel Interface—**tunnel.2**

IP Pool—**10.31.32.3 - 10.31.32.118**

STEP 7 | Configure the GlobalProtect Portals.

© 2017 Palo Alto Networks, Inc.

Select **Network** > **GlobalProtect** > **Portals** and add the following configuration:

1. Set Up Access to the GlobalProtect Portal:

   Interface—`ethernet1/2`

   IP Address—`203.0.113.1`

   Server Certificate—`GP-server-cert.pem issued by GoDaddy`

   Certificate Profile—`GP-client-cert`
2. Define the GlobalProtect Agent Configurations:

   Connect Method—`On-demand` (Manual user initiated connection)

   External Gateway Address—`gp.acme.com`

STEP 8 | Deploy the GlobalProtect Agent Software.

Select **Device** > **GlobalProtect Client**.

In this example, use the procedure to Host Agent Updates on the Portal.

STEP 9 | (**Optional**) Enable use of the GlobalProtect mobile app.

Purchase and install a GlobalProtect subscription (**Device** > **Licenses**) to enable use of the app.

STEP 10 | Save the GlobalProtect configuration.

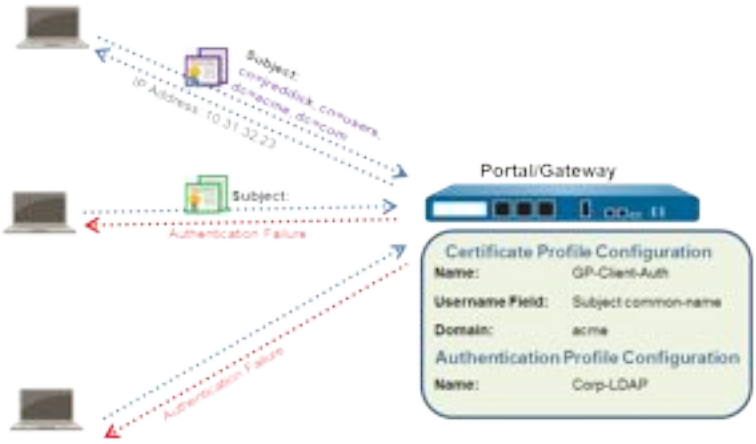
Click **Commit**.

© 2017 Palo Alto Networks, Inc.

Remote Access VPN with Two-Factor Authentication

If you configure a GlobalProtect portal or gateway with an authentication profile and a certificate profile (which together can provide two-factor authentication), the end user must succeed at authentication through both profiles before gaining access. For portal authentication, this means that certificates must be pre-deployed to the end clients before their initial portal connection. Additionally, the client certificate presented by a client must match what is defined in the certificate profile.

- If the certificate profile does not specify a username field (that is, the **Username Field** it is set to **None**), the client certificate does not need to have a username. In this case, the client must provide the username when authenticating against the authentication profile.
- If the certificate profile specifies a username field, the certificate that the client presents must contain a username in the corresponding field. For example, if the certificate profile specifies that the username field is Subject, the certificate presented by the client must contain a value in the common-name field, or else the authentication fails. In addition, when the username field is required, the value from the username field of the certificate is automatically populated as the username when the user attempts to enter credentials for authenticating to the authentication profile. If you do not want force users to authenticate with a username from the certificate, do not specify a username field in the certificate profile.



This quick configuration uses the same topology as GlobalProtect VPN for Remote Access. However, in this configuration the clients must authenticate against a certificate profile and an authentication profile. For more details on a specific type of two-factor authentication, see the following topics:

- Enable Two-Factor Authentication Using Certificate and Authentication Profiles
- Enable Two-Factor Authentication Using One-Time Passwords (OTPs)
- Enable Two-Factor Authentication Using Smart Cards

Use the following procedure to configure VPN Remote Access with Two-Factor Authentication.

STEP 1 | Create Interfaces and Zones for GlobalProtect.

Use the default virtual router for all interface configurations to avoid having to create inter-zone routing.

© 2017 Palo Alto Networks, Inc.

- Select **Network** > **Interfaces** > **Ethernet** and configure ethernet1/2 as a Layer 3 Ethernet interface with IP address **`203.0.113.1`** and assign it to the l3-untrust security zone and the default virtual router.
- Create a DNS "A" record that maps IP address **`203.0.113.1`** to **`gp.acme.com`**.
- Select **Network** > **Interfaces** > **Tunnel** and add the tunnel.2 interface and add it to a new zone called **`corp-vpn`**. Assign it to the default virtual router.
- Enable User Identification on the corp-vpn zone.

STEP 2 | Create security policy to enable traffic flow between the corp-vpn zone and the l3-trust zone to enable access to your internal resources.

1. Select **Policies** > **Security** and then click **Add** to add a new rule.
2. For this example, you would define the rule with the following settings:
   - Name—**`VPN Access`**
   - Source Zone—**`corp-vpn`**
   - Destination Zone—**`l3-trust`**



STEP 3 | Obtain a server certificate for the interface hosting the GlobalProtect portal and gateway using one of the following methods:

- (**Recommended**) Import a server certificate from a well-known, third-party CA.
- Use the root CA on the portal to generate a self-signed server certificate.

Select **Device** > **Certificate Management** > **Certificates** to manage certificates as follows:

- Obtain a server certificate. Because the portal and gateway are on the same interface, the same server certificate can be used for both components.
- The CN of the certificate must match the FQDN, **`gp.acme.com`**.
- To enable clients to connect to the portal without receiving certificate errors, use a server certificate from a public CA.

STEP 4 | Issue client certificates to GlobalProtect clients and endpoints.

1. Use your enterprise PKI or a public CA to issue a unique client certificate to each GlobalProtect user.
2. Install certificates in the personal certificate store on the endpoints.

STEP 5 | Create a client certificate profile.

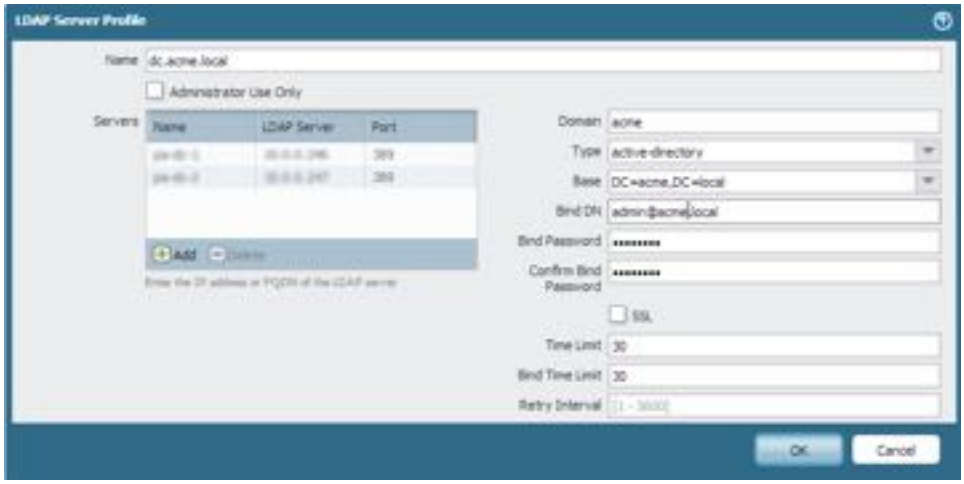
1. Select **Device** > **Certificate Management** > **Certificate Profile**, **Add** and enter a profile **Name** such as **`GP-client-cert`**.
2. Specify where to get the username that will be used to authenticate the end user:
   - **From user**—If you want the end user to supply a username when authenticating to the service specified in the authentication profile, select **None** as the **Username Field**.
   - **From certificate**—If you want to extract the username from the certificate, select **Subject** as the **Username Field**. If you use this option, the CN contained in the certificate will automatically populated the username field when the user is prompted to login to the portal/gateway and the user will be required to log in using that username.
3. In the CA Certificates section, **Add** and then select the **CA Certificate** that issued the client certificates, and click **OK** twice.

© 2017 Palo Alto Networks, Inc.

STEP 6 | Create a server profile.

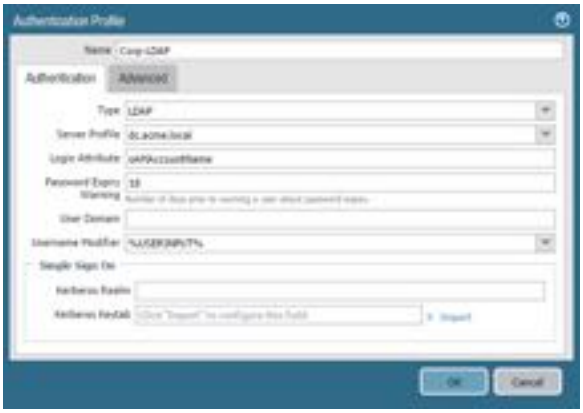
The server profile instructs the firewall how to connect to the authentication service. Local, RADIUS, Kerberos, SAML, and LDAP authentication methods are supported. This example shows an LDAP authentication profile for authenticating users against the Active Directory.

Create the server profile for connecting to the LDAP server (**Device** > **Server Profiles** > **LDAP**).



STEP 7 | (Optional) Create an authentication profile.

Attach the server profile to an authentication profile (**Device > Authentication Profile**).



STEP 8 | Configure a GlobalProtect Gateway.

See the topology diagram shown in GlobalProtect VPN for Remote Access.

Select **Network** > **GlobalProtect** > **Gateways** and add the following configuration:

Interface—**`ethernet1/2`**

IP Address—**`203.0.113.1`**

Server Certificate—**`GP-server-cert.pem issued by GoDaddy`**

Certificate Profile—**`GP-client-cert`**

Authentication Profile—**`Corp-LDAP`**

Tunnel Interface—**`tunnel.2`**

IP Pool—**`10.31.32.3 - 10.31.32.118`**

© 2017 Palo Alto Networks, Inc.

STEP 9 | Configure the GlobalProtect Portals.

Select **Network** > **GlobalProtect** > **Portals** and add the following configuration:

1. Set Up Access to the GlobalProtect Portal:

   Interface—**`ethernet1/2`**

   IP Address—**`203.0.113.1`**

   Server Certificate—**`GP-server-cert.pem issued by GoDaddy`**

   Certificate Profile—**`GP-client-cert`**

   Authentication Profile—**`Corp-LDAP`**
2. Define the GlobalProtect Agent Configurations:

   Connect Method—**`On-demand`** (Manual user initiated connection)

   External Gateway Address—**`gp.acme.com`**

STEP 10 | Deploy the GlobalProtect Agent Software.

Select **Device** > **GlobalProtect Client.**

In this example, use the procedure to Host Agent Updates on the Portal.

STEP 11 | (**Optional**) Deploy Agent Settings Transparently.

As an alternative to deploying agent settings from the portal configuration, you can define settings directly from the Windows registry or global MAC plist. Examples of settings that you can deploy include specifying the portal IP address or enabling GlobalProtect to initiate a VPN tunnel before a user logs in to the device and connects to the GlobalProtect portal. On Windows clients only, you can also configure settings using the MSIEXEC installer. For additional information, see Customizable Agent Settings.

STEP 12 | (**Optional**) Enable use of the GlobalProtect mobile app.

Purchase and install a GlobalProtect subscription (**Device** > **Licenses**) to enable use of the app.

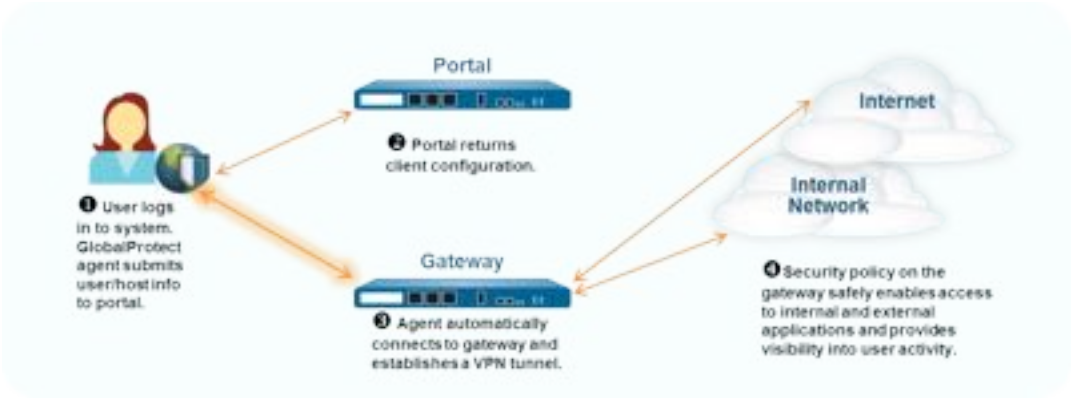
STEP 13 | Save the GlobalProtect configuration.

Click **Commit**.

© 2017 Palo Alto Networks, Inc.

Always On VPN Configuration

In an “always on” GlobalProtect configuration, the agent connects to the GlobalProtect portal upon user logon to submit user and host information and receive the client configuration. It then automatically establishes the VPN tunnel to the gateway specified in the client configuration delivered by the portal without end user intervention as shown in the following illustration.



To switch any of the previous remote access VPN configurations to an always-on configuration, you simply change the connect method:

- Remote Access VPN (Authentication Profile) on page 201
- Remote Access VPN (Certificate Profile) on page 204
- Remote Access VPN with Two-Factor Authentication on page 207

Use the following procedure to switch to an *Always On* configuration.

STEP 1 | Select **Network** > **GlobalProtect** > **Portals** and select the portal configuration to open it.

STEP 2 | Select the **Agent** tab and then select the agent configuration you want to modify.

STEP 3 | Select the **App** tab.

STEP 4 | Select **User-logon (Always On)** as the **Connect Method**. Repeat this step for each agent configuration.

STEP 5 | Click **OK** twice to save the agent configuration and the portal configuration and then **Commit** your changes.

© 2017 Palo Alto Networks, Inc.

Remote Access VPN with Pre-Logon

Pre-logon is a connect method that establishes a VPN tunnel before a user logs in. The purpose of pre-logon is to authenticate the endpoint (not the user) and then enable domain scripts and other tasks of your choice to run as soon as the endpoint powers on. A machine certificate enables the endpoint to have the VPN tunnel to the gateway. A common practice for IT personnel is to install the machine certificate while staging the endpoint for the user.

A pre-logon VPN tunnel has no username association because the user has not logged in. Therefore, to let the endpoint have access to resources in the trust zone, you must create security policies that match the pre-logon user. These policies should allow access to only the basic services for starting up the system, such as DHCP, DNS, Active Directory (for example, to change an expired password), antivirus, or operating system update services.

After the gateway authenticates a Windows user, the VPN tunnel is reassigned to that user (the IP address mapping on the firewall changes from the pre-logon endpoint to the authenticated user).



Mac systems behave differently from Windows systems with pre-logon. With Mac OS, the tunnel created for pre-logon is torn down and a new tunnel created when the user logs in.

When a client requests a new connection, the portal authenticates the client by using an authentication profile. The portal can also use an optional certificate profile that validates the client certificate (if the configuration includes a client certificate). In this case, the client certificate must identify the user.

After authentication, the portal determines if the client's configuration is current. If the portal's configuration for the agent has changed, it pushes an updated configuration to the endpoint.

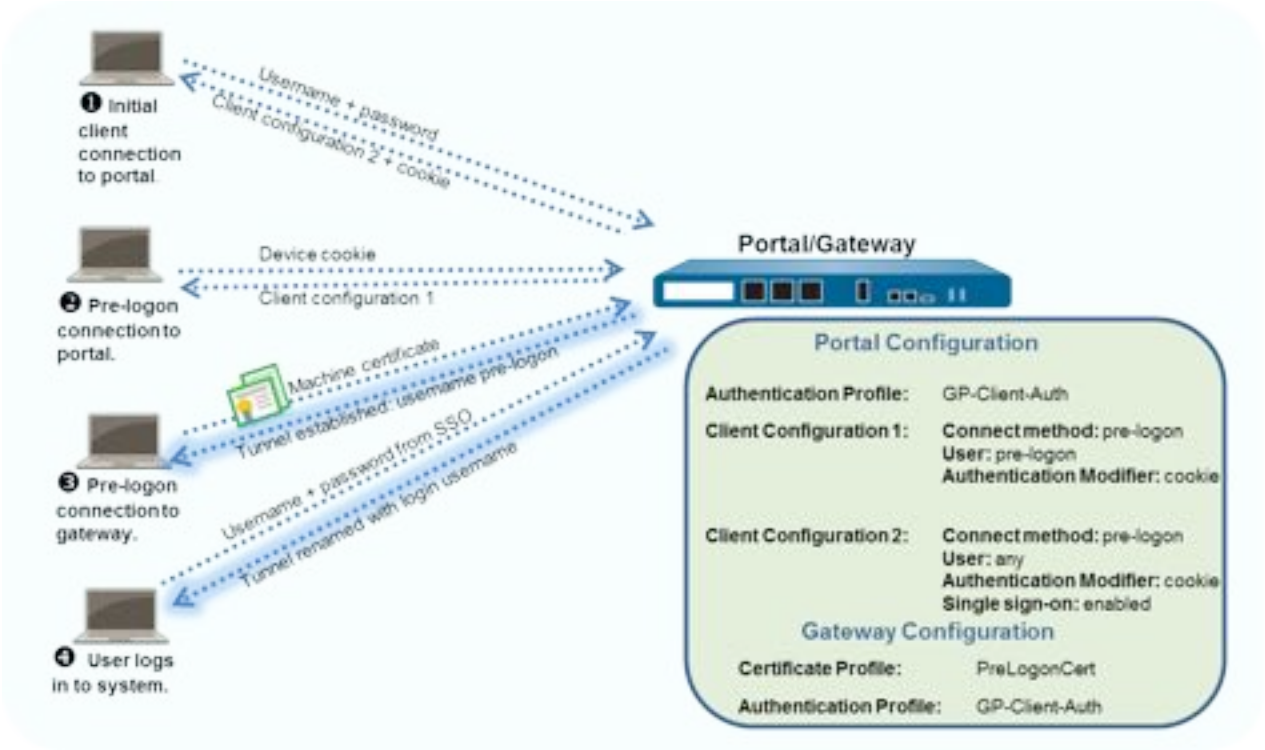
If the configuration on the portal or a gateway includes cookie-based authentication for the client, the portal or gateway installs an encrypted cookie on the client. Subsequently, the portal or gateway uses the cookie to authenticate users and for refreshing the client's configuration. Also, if an agent configuration profile includes the pre-logon connect method in addition to cookie-authentication, the GlobalProtect components can use the cookie for pre-logon.

If users never log into a device (for example, a headless device) or a pre-logon connection is required on a system that a user has not previously logged into, you can let the endpoint initiate a pre-logon tunnel without first connecting to the portal to download the pre-logon configuration. To do this, you must override the default behavior by creating entries in the Windows registry or Mac plist.

The GlobalProtect endpoint will then connect to the portal specified in the configuration and authenticate the endpoint by using its machine certificate (as specified in a certificate profile configured on the gateway) and establish the VPN tunnel.

When the end user subsequently logs in to the machine and if single sign-on (SSO) is enabled in the client configuration, the username and password are captured while the user logs in and used to authenticate to the gateway and so that the tunnel can be renamed (Windows). If SSO is not enabled in the client configuration or of SSO is not supported on the client system (for example, it is a Mac OS system) the users' credentials must be stored in the agent (that is, the **Save User Credentials** option must be set to **Yes**). After successful authentication to the gateway the tunnel will be renamed (Windows) or rebuilt (Mac) and user- and group-based policy can be enforced.

© 2017 Palo Alto Networks, Inc.



This example uses the GlobalProtect topology shown in GlobalProtect VPN for Remote Access.

STEP 1 | Create Interfaces and Zones for GlobalProtect.

Use the default virtual router for all interface configurations to avoid having to create inter-zone routing.

- For this example, select **Network** > **Interfaces** > **Ethernet** and then:
- Select **ethernet1/2**.
- For its interface type, select **Layer 3**.
- **Assign interface to**: default virtual router, default virtual system, and **l3-untrust** security zone.
- Select **IPv4** and **Add**.
- Select the address **`203.0.113.1`** (or the object that maps **`203.0.113.1`**) or add a **New Address** to create a new object and address mapping. (Leave the address type as **Static**.)
- Create a DNS "A" record that maps IP address **`203.0.113.1`** to **`gp.acme.com`**.
- Select **Network** > **Interfaces** > **Tunnel**.
- **Add** a tunnel.2 interface to a new zone called **`corp-vpn`**. Assign it to the default virtual router.
- Enable User Identification on the corp-vpn zone.

STEP 2 | Create the security policy rules.

This configuration requires the following policies (**Policies** > **Security**):

1. Create a rule that enables pre-logon user access to basic services that are required for the computer to come up, such as authentication services, DNS, DHCP, and Microsoft Updates.
2. Create a rule to deny pre-logon user access to all other destinations and applications.
3. Create any additional rules to enable access to specific destinations and applications for specific users or user groups. Follow the Best Practice Internet Gateway Security Policy recommendations for creating these rules.

© 2017 Palo Alto Networks, Inc.

STEP 3 | Use one of the following methods to obtain a server certificate for the interface that is hosts the GlobalProtect portal and gateway:

- (**Recommended**) Import a server certificate from a well-known, third-party CA.
- Use the root CA on the portal to generate a self-signed server certificate.

Select **Device** > **Certificate Management** > **Certificates** to manage certificates with the following criteria:

- Obtain a server certificate. Because the portal and gateway are on the same interface, the same server certificate can be used for both components.
- The CN of the certificate must match the FQDN, **`gp.acme.com`**.
- To enable clients to connect to the portal without receiving certificate errors, use a server certificate from a public CA.

STEP 4 | Generate a machine certificate for each client system that will connect to GlobalProtect and import them into the personal certificate store on each machine.

Although you could generate self-signed certificates for each client system, as a best practice, use your own public-key infrastructure (PKI) to issue and distribute certificates to your clients.

1. Issue client certificates to GlobalProtect clients and endpoints.
2. Install certificates in the personal certificate store on the endpoints. (Local Computer store on Windows or System Keychain on Mac OS)

STEP 5 | Import the trusted root CA certificate from the CA that issued the machine certificates onto the portal and gateway(s).



You do not have to import the private key.

1. Download the CA certificate in Base64 format.
2. Import the certificate onto each firewall that hosts a portal or gateway, as follows:
    1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
    2. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.
    3. **Browse** to the **Certificate File** you downloaded from the CA.
    4. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.
    5. Select the certificate you just imported on the **Device Certificates** tab to open it.
    6. Select **Trusted Root CA** and then click **OK**.

STEP 6 | On each firewall that hosts a GlobalProtect gateway, create a certificate profile to identify the CA certificate for validating the machine certificates.

Optionally, if you plan to use client certificate authentication to authenticate users when they log in to the system, make sure that the CA certificate that issues the client certificates is referenced in the certificate profile in addition to the CA certificate that issued the machine certificates if they are different.

1. Select **Device** > **Certificates** > **Certificate Management** > **Certificate Profile**.
2. Click **Add** and enter a **Name** to uniquely identify the profile, such as **`PreLogonCert`**.
3. Set **Username Field** to **None**.
4. (**Optional**) If you will also use client certificate authentication to authenticate users upon login, add the CA certificate that issued the client certificates if it is different from the one that issued the machine certificates.
5. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you imported in 5 and then click **OK**.

© 2017 Palo Alto Networks, Inc.

6. Click **OK** to save the profile.

STEP 7 | Configure a GlobalProtect Gateway.

See the topology diagram shown in GlobalProtect VPN for Remote Access.

Although you must create a certificate profile for pre-logon access to the gateway, you can use either client certificate authentication or authentication profile-based authentication for logged in users. In this example, the same LDAP profile is used that is used to authenticate users to the portal.

1. Select **Network** > **GlobalProtect** > **Gateways** and add the following configuration:

   Interface—**`ethernet1/2`**

   IP Address—**`203.0.113.1`**

   Server Certificate—**`GP-server-cert.pem issued by GoDaddy`**

   Certificate Profile—**`PreLogonCert`**

   Authentication Profile—**`Corp-LDAP`**

   Tunnel Interface—**`tunnel.2`**

   IP Pool—**`10.31.32.3 - 10.31.32.118`**
2. **Commit** the gateway configuration.

STEP 8 | Configure the GlobalProtect Portals.

Configure **Device** details (networking parameters, the authentication service profile, and the certificate for the authentication server).

Select **Network** > **GlobalProtect** > **Portals** and specify the following configuration:

Set Up Access to the GlobalProtect Portal:

Interface—**`ethernet1/2`**

IP Address—**`203.0.113.1`**

Server Certificate—**`GP-server-cert.pem issued by GoDaddy`**

Certificate Profile—**`None`**

Authentication Profile—`Corp-LDAP`

STEP 9 | Define the GlobalProtect Agent Configurations for pre-logon users and for logged in users.

Use a single agent configuration if you want pre-logon users to access the same gateways before and after they log in.

Otherwise, to direct pre-logon users to different gateways before and after they log in, create two agent configuration profiles. In this first agent configuration's **User/User Group**, select the **pre-logon** filter. With pre-logon, the portal first authenticates the endpoint, not the user, to set up a VPN (even though the pre-logon parameter is associated with users). Subsequently, the portal authenticates the user when he or she logs in.

After the portal authenticates the user, it deploys the second agent configuration. In this case, **User/User Group** is **any**.



As a best practice, enable SSO in the second agent configuration so that the correct username is immediately reported to the gateway when the user logs in to the endpoint. If SSO is not enabled, the saved username in the Agent settings panel is used.

Select **Agent** and specify one of the following configurations:

© 2017 Palo Alto Networks, Inc.

- Use the same gateway before and after pre-logon users log in:

Use single sign-on—**`enabled`**

Connect Method—**`pre-logon`**

External Gateway Address—**`gp1.acme.com`**

User/User Group—**`any`**

Authentication Override—Cookie authentication for transparently authenticating users and for configuration refresh

- Use separate gateways for pre-logon users before and after they log in:

First Agent Configuration:

Connect Method—**`pre-logon`**

External Gateway Address—**`gp1.acme.com`**

User/User Group—**`pre-logon`**

Authentication Override—Cookie authentication for transparently authenticating users and for configuration refresh

Second Agent Configuration:

Use single sign-on—**`enabled`**

Connect Method—**`pre-logon`**

External Gateway Address—**`gp2.acme.com`**

User/User Group—**`any`**

Authentication Override—Cookie authentication for transparently authenticating users and for configuration refresh

Make sure the pre-logon client configuration is first in the list of configurations. If it is not, select it and click **Move Up**.

STEP 10 | Save the GlobalProtect configuration.

Click **Commit**.

STEP 11 | (**Optional**) If users will never log into a device (for example, a headless device) or a pre-logon connection is required on a system that a user has not previously logged into, create the **`Prelogon`** registry entry on the client system.

You must also pre-deploy additional agent settings such as the default portal IP address and connect method.

For more information about registry settings, see Deploy Agent Settings Transparently.

1. Locate the GlobalProtect settings in the registry:

   `HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\PanSetup`
2. Create a **DWORD** named **`Prelogon`** with a value of **`1`** in the **Value data** field and **Hexadecimal** as the **Base**. This setting enables GlobalProtect to initiate a VPN connection before the user logs into the laptop.
3. Create a **String Value** named **`Portal`** that specifies the IP address or hostname of the default portal for the GlobalProtect client.

© 2017 Palo Alto Networks, Inc.

4. Create a **String Value** named **`connect-method`** with a value of **`pre-logon`** in the Value data field. This setting enables GlobalProtect to initiate a VPN tunnel before a user logs in to the device and connects to the GlobalProtect portal.

© 2017 Palo Alto Networks, Inc.

GlobalProtect Multiple Gateway Configuration

In GlobalProtect Multiple Gateway Topology, a second external gateway has been added to the configuration. Multiple gateways are supported in all of the preceding example configurations. Additional steps include configuring a second firewall as a GlobalProtect gateway. In addition, when configuring the client configurations to be deployed by the portal you can decide whether to allow access to all gateways, or specify different gateways for different configurations.

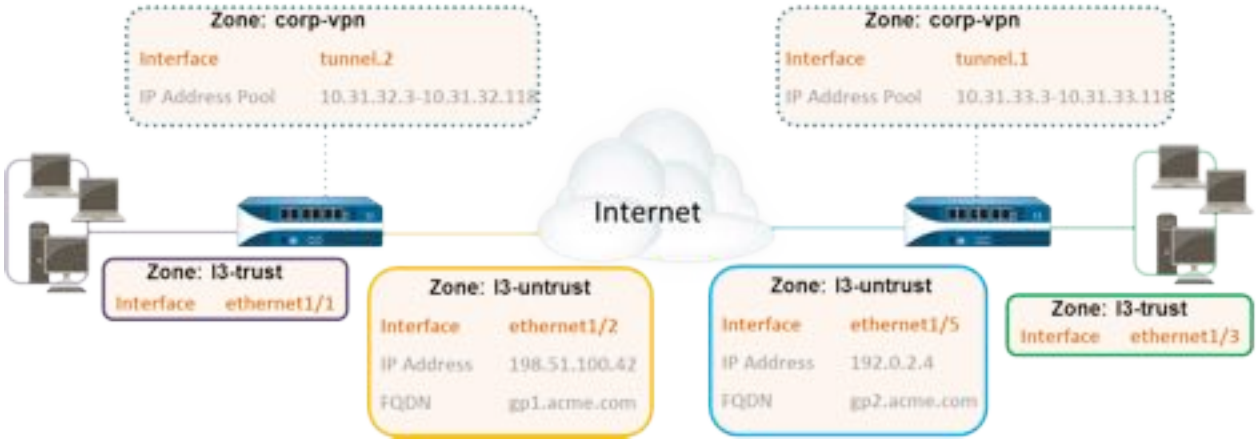


Figure 5: GlobalProtect Multiple Gateway Topology

If a client configuration contains more than one gateway, the agent will attempt to connect to all gateways listed in its client configuration. The agent will then use priority and response time as to determine the gateway to which to connect. The agent connects to a lower priority gateway only if the response time for the higher priority gateway is greater than the average response time across all gateways. For more information, see Gateway Priority in a Multiple Gateway Configuration.

STEP 1 | Create Interfaces and Zones for GlobalProtect.

In this configuration, you must set up interfaces on each firewall hosting a gateway.

Use the default virtual router for all interface configurations to avoid having to create inter-zone routing.

On the firewall hosting the portal/gateway (gw1):

- Select **Network** > **Interfaces** > **Ethernet** and configure ethernet1/2 as a Layer 3 Ethernet interface with IP address **198.51.100.42** and assign it to the l3-untrust security zone and the default virtual router.
- Create a DNS "A" record that maps IP address **198.51.100.42** to **gp1.acme.com**.
- Select **Network** > **Interfaces** > **Tunnel** and add the tunnel.2 interface and add it to a new zone called `corp-vpn`. Assign it to the default virtual router.
- Enable User Identification on the corp-vpn zone.

On the firewall hosting the second gateway (gw2):

- Select **Network** > **Interfaces** > **Ethernet** and configure ethernet1/5 as a Layer 3 Ethernet interface with IP address **192.0.2.4** and assign it to the l3-untrust security zone and the default virtual router.
- Create a DNS "A" record that maps IP address **192.0.2.4** to **gp2.acme.com**.

© 2017 Palo Alto Networks, Inc.

- Select **Network** > **Interfaces** > **Tunnel** and add the tunnel.1 interface and add it to a new zone called `corp-vpn`. Assign it to the default virtual router.
- Enable User Identification on the corp-vpn zone.

STEP 2 | Purchase and install a GlobalProtect subscription on each gateway if you have users who will be using the GlobalProtect app on their mobile devices or if you plan to use HIP-enabled security policy.

After you purchase the GlobalProtect subscription and receive your activation code, install the license on the firewall hosting the portal as follows:

1. Select **Device** > **Licenses**.
2. Select **Activate feature using authorization code**.
3. When prompted, enter the **Authorization Code** and then click **OK**.
4. Verify that the license was successfully activated.



STEP 3 | On each firewall hosting a GlobalProtect gateway, create security policy.

This configuration requires policy rules to enable traffic flow between the corp-vpn zone and the l3-trust zone to enable access to your internal resources (**Policies** > **Security**).



STEP 4 | Obtain server certificates for the interfaces hosting your GlobalProtect portal and each of your GlobalProtect gateways using the following recommendations:

- (**On the firewall hosting the portal or portal/gateway**) Import a server certificate from a well-known, third-party CA.
- (**On a firewall hosting only a gateway**) Use the root CA on the portal to generate a self-signed server certificate.

On each firewall hosting a portal/gateway or gateway, select **Device** > **Certificate Management** > **Certificates** to manage certificates as follows:

- Obtain a server certificate for the portal/gw1. Because the portal and the gateway are on the same interface you must use the same server certificate. The CN of the certificate must match the FQDN, gp1.acme.com. To enable clients to connect to the portal without receiving certificate errors, use a server certificate from a public CA.
- Obtain a server certificate for the interface hosting gw2. Because this interface hosts a gateway only you can use a self-signed certificate. The CN of the certificate must match the FQDN, gp2.acme.com.

STEP 5 | Define how you will authenticate users to the portal and the gateways.

You can use any combination of certificate profiles and/or authentication profiles as necessary to ensure the security for your portal and gateways. Portals and individual gateways can also use different authentication schemes. See the following sections for step-by-step instructions:

- Set Up External Authentication (authentication profile)
- Set Up Client Certificate Authentication (certificate profile)
- Set Up Two-Factor Authentication (token- or OTP-based)

© 2017 Palo Alto Networks, Inc.

You will then need to reference the certificate profile and/or authentication profiles you defined in the portal and gateway configurations you define.

STEP 6 | Configure the gateways.

This example shows the configuration for gp1 and gp2 shown in GlobalProtect Multiple Gateway Topology. (See Configure a GlobalProtect Gateway for step-by-step instructions on creating the gateway configurations.)

On the firewall hosting gp1, select **Network** > **GlobalProtect** > **Gateways** and configure the gateway settings as follows:

Interface—**`ethernet1/2`**

IP Address—**`198.51.100.42`**

Server Certificate—**`GP1-server-cert.pem issued by GoDaddy`**

Tunnel Interface—**`tunnel.2`**

IP Pool—**`10.31.32.3 - 10.31.32.118`**

On the firewall hosting gp2, select **Network** > **GlobalProtect** > **Gateways** and configure the gateway settings as follows:

Interface—**`ethernet1/2`**

IP Address—**`192.0.2.4`**

Server Certificate—**`self-signed certificate, GP2-server-cert.pem`**

Tunnel Interface—**`tunnel.1`**

IP Pool—**`10.31.33.3 - 10.31.33.118`**

STEP 7 | Configure the GlobalProtect Portals.

Select **Network** > **GlobalProtect** > **Portals** and add the following configuration:

1. Set Up Access to the GlobalProtect Portal:

   Interface—**`ethernet1/2`**

   IP Address—**`198.51.100.42`**

   Server Certificate—**`GP1-server-cert.pem issued by GoDaddy`**
2. Define the GlobalProtect Agent Configurations:

   The number of client configurations you create depends on your specific access requirements, including whether you require user/group-based policy and/or HIP-enabled policy enforcement.

STEP 8 | Deploy the GlobalProtect Agent Software.

Select **Device** > **GlobalProtect Client**.

In this example, use the procedure to Host Agent Updates on the Portal.

STEP 9 | Save the GlobalProtect configuration.

Click **Commit** on the firewall hosting the portal and the gateway(s).

© 2017 Palo Alto Networks, Inc.

GlobalProtect for Internal HIP Checking and User-Based Access

When used in conjunction with User-ID and/or HIP checks, an internal gateway can be used to provide a secure, accurate method of identifying and controlling traffic by user and/or device state, replacing other network access control (NAC) services. Internal gateways are useful in sensitive environments where authenticated access to critical resources is required.

In a configuration with only internal gateways, all clients must be configured with user-logon; on-demand mode is not supported. In addition, it is recommended that you configure all client configurations to use single sign-on (SSO). Additionally, because internal hosts do not need to establish a tunnel connection with the gateway, the IP address of the physical network adapter on the client system is used.

In this quick config, internal gateways are used to enforce group based policies that allow users in the Engineering group access to the internal source control and bug databases and users in the Finance group to the CRM applications. All authenticated users have access to internal web resources. In addition, HIP profiles configured on the gateway check each host to ensure compliance with internal maintenance requirements, such as whether the latest security patches and antivirus definitions are installed, whether disk encryption is enabled, or whether the required software is installed.

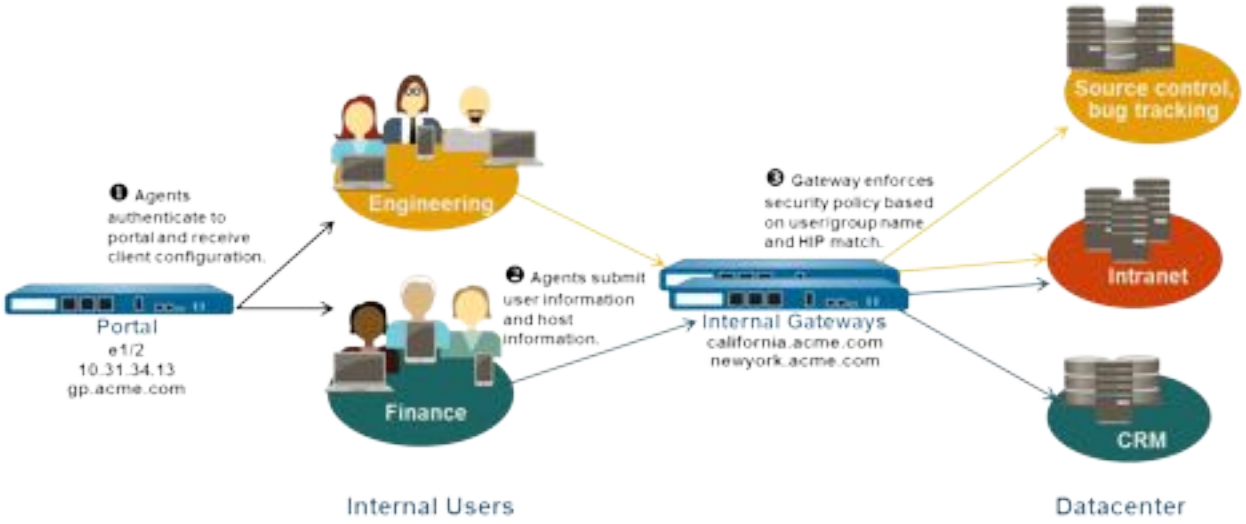


Figure 6: GlobalProtect Internal Gateway Configuration

Use the following procedure to quickly configure a GlobalProtect internal gateway.

STEP 1 | Create Interfaces and Zones for GlobalProtect.

In this configuration, you must set up interfaces on each firewall hosting a portal and/or a gateway. Because this configuration uses internal gateways only, you must configure the portal and gateways on interfaces on the internal network.

Use the default virtual router for all interface configurations to avoid having to create inter-zone routing.

On each firewall hosting a portal/gateway:

© 2017 Palo Alto Networks, Inc.

1. Select an Ethernet port to host the portal/gateway and then configure a Layer3 interface with an IP address in the l3-trust security zone. (**Network** > **Interfaces** > **Ethernet**).
2. **Enable User Identification** on the l3-trust zone.

STEP 2 | Purchase and install a GlobalProtect subscription for each firewall hosting an internal gateway if you have users who will be using the GlobalProtect app on their mobile devices or if you plan to use HIP-enabled security policy.



After you purchase the GlobalProtect subscriptions and receive your activation code, install the GlobalProtect subscriptions on the firewalls hosting your gateways as follows:

1. Select **Device** > **Licenses**.
2. Select **Activate feature using authorization code**.
3. When prompted, enter the **Authorization Code** and then click **OK**.
4. Verify that the license was successfully activated.

Contact your Palo Alto Networks Sales Engineer or Reseller if you do not have the required licenses. For more information on licensing, see About GlobalProtect Licenses.

STEP 3 | Obtain server certificates for the GlobalProtect portal and each GlobalProtect gateway.

In order to connect to the portal for the first time, the end clients must trust the root CA certificate used to issue the portal server certificate. You can either use a self-signed certificate on the portal and deploy the root CA certificate to the end clients before the first portal connection, or obtain a server certificate for the portal from a trusted CA.

You can use self-signed certificates on the gateways.

The recommended workflow is as follows:

1. On the firewall hosting the portal:
    1. Import a server certificate from a well-known, third-party CA.
    2. Create the root CA certificate for issuing self-signed certificates for the GlobalProtect components.
    3. Use the root CA on the portal to generate a self-signed server certificate. Repeat this step for each gateway.
2. On each firewall hosting an internal gateway: Deploy the self-signed server certificates.

STEP 4 | Define how you will authenticate users to the portal and the gateways.

You can use any combination of certificate profiles and/or authentication profiles as necessary to ensure the security for your portal and gateways. Portals and individual gateways can also use different authentication schemes. See the following sections for step-by-step instructions:

- Set Up External Authentication (authentication profile)
- Set Up Client Certificate Authentication (certificate profile)
- Set Up Two-Factor Authentication (token- or OTP-based)

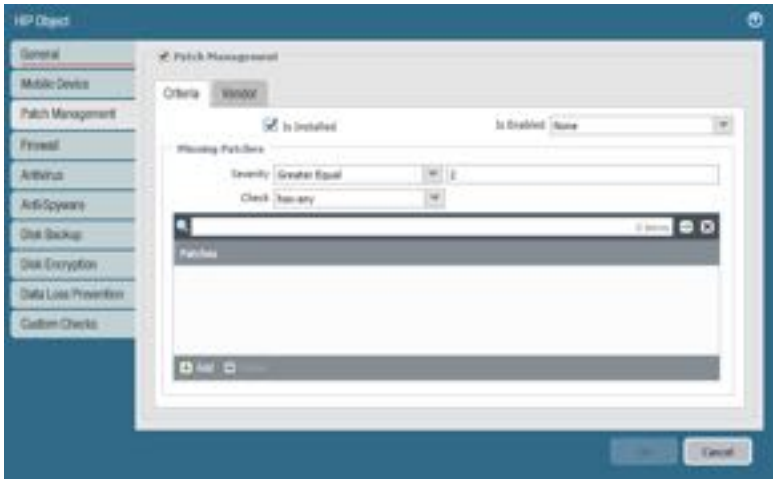
You will then need to reference the certificate profile and/or authentication profiles you defined in the portal and gateway configurations you define.

STEP 5 | Create the HIP profiles you will need to enforce security policy on gateway access.

See Host Information for more information on HIP matching.

© 2017 Palo Alto Networks, Inc.

1. Create the HIP objects to filter the raw host data collected by the agents. For example, if you are interested in preventing users that are not up to date with required patches, you might create a HIP object to match on whether the patch management software is installed and that all patches with a given severity are up to date.



2. Create the HIP profiles that you plan to use in your policies.

   For example, if you want to ensure that only Windows users with up-to-date patches can access your internal applications, you might attach the following HIP profile that will match hosts that do NOT have a missing patch:



STEP 6 | Configure the internal gateways.

Select **Network** > **GlobalProtect** > **Gateways** and add the following settings:

- **Interface**
- **IP Address**
- **Server Certificate**
- **Authentication Profile** and/or **Configuration Profile**

Notice that it is not necessary to configure the client configuration settings in the gateway configurations (unless you want to set up HIP notifications) because tunnel connections are not required. See Configure a GlobalProtect Gateway for step-by-step instructions on creating the gateway configurations.

STEP 7 | Configure the GlobalProtect Portals.

Although all of the previous configurations could use a Connect Method of User-logon (Always On) or On-demand (Manual user initiated connection), an internal gateway configuration must always be on and therefore requires a Connect Method of User-logon (Always On).

© 2017 Palo Alto Networks, Inc.

Select **Network** > **GlobalProtect** > **Portals** and add the following configuration:

1. Set Up Access to the GlobalProtect Portal:

   Interface—**`ethernet1/2`**

   IP Address—**`10.31.34.13`**

   Server Certificate—**`GP-server-cert.pem issued by GoDaddy`** with **`CN=gp.acme.com`**
2. Define the GlobalProtect Client Authentication Configurations:

   Use single sign-on—**`enabled`**

   Connect Method—**`User-logon (Always On)`**

   Internal Gateway Address—**`california.acme.com, newyork.acme.com`**

   User/User Group—**`any`**
3. **Commit** the portal configuration.

STEP 8 | Deploy the GlobalProtect Agent Software.

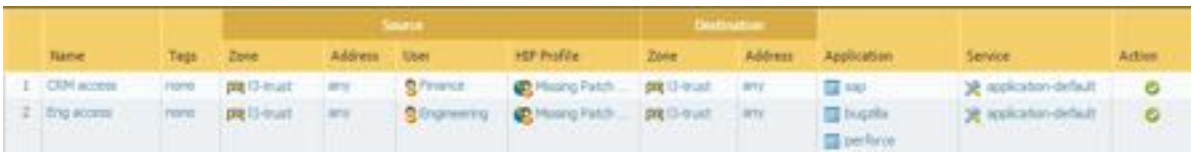
Select **Device** > **GlobalProtect Client**.

In this example, use the procedure to Host Agent Updates on the Portal.

STEP 9 | Create the HIP-enabled and/or user/group-based security rules on your gateway(s).

Add the following security rules for this example:

1. Select **Policies** > **Security** and click **Add**.
2. On the **Source** tab, set the **Source Zone** to **l3-trust**.
3. On the **User** tab, add the HIP profile and user/group to match.
   - Click **Add** in the **HIP Profiles** section and select the HIP profile **MissingPatch**.
   - Click **Add** in the **Source User** section and select the group (Finance or Engineering depending on which rule you are creating).
4. Click **OK** to save the rule.
5. **Commit** the gateway configuration.



© 2017 Palo Alto Networks, Inc.

Mixed Internal and External Gateway Configuration

In a GlobalProtect mixed internal and external gateway configuration, you configure separate gateways for VPN access and for access to your sensitive internal resources. With this configuration, agents perform internal host detection to determine if they are on the internal or external network. If the agent determines it is on the external network, it will attempt to connect to the external gateways listed in its client configuration and it will establish a VPN (tunnel) connection with the gateway with the highest priority and the shortest response time.

Because security policies are defined separately on each gateway, you have granular control over which resources your external and internal users have access to. In addition, you also have granular control over which gateways users have access to by configuring the portal to deploy different client configurations based on user/group membership or based on HIP profile matching.

In this example, the portals and all three gateways (one external and two internal) are deployed on separate firewalls. The external gateway at gpvpn.acme.com provides remote VPN access to the corporate network while the internal gateways provide granular access to sensitive datacenter resources based on group membership. In addition, HIP checks are used to ensure that hosts accessing the datacenter are up-to-date on security patches.

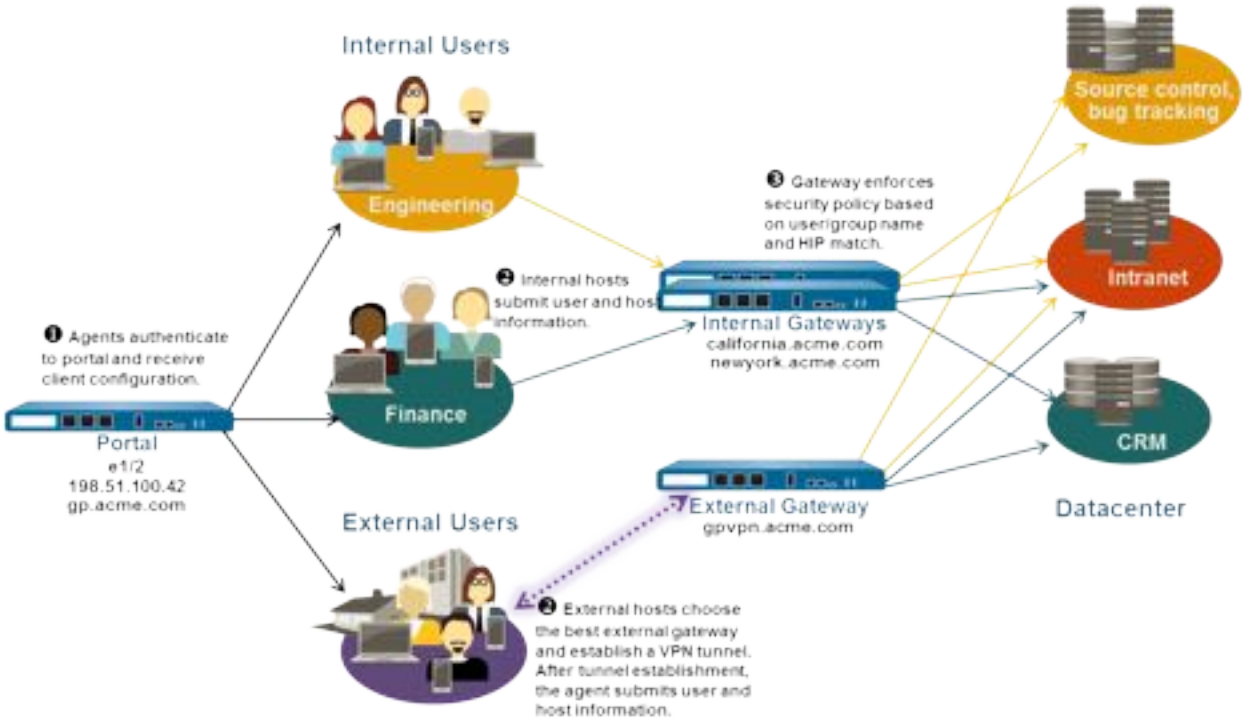


Figure 7: GlobalProtect Deployment with Internal and External Gateways

Use the following procedure to quickly configure a mix of internal and external GlobalProtect gateways.

STEP 1 | Create Interfaces and Zones for GlobalProtect.

In this configuration, you must set up interfaces on the firewall hosting a portal and each firewall hosting a gateway.

© 2017 Palo Alto Networks, Inc.

Use the default virtual router for all interface configurations to avoid having to create inter-zone routing.

On the firewall hosting the portal gateway (gp.acme.com):

- Select **Network** > **Interfaces** > **Ethernet** and configure ethernet1/2 as a Layer 3 Ethernet interface with IP address **198.51.100.42** and assign it to the l3-untrust security zone and the default virtual router.
- Create a DNS "A" record that maps IP address 198.51.100.42 to gp.acme.com.
- Select **Network** > **Interfaces** > **Tunnel** and add the tunnel.2 interface and add it to a new zone called `corp-vpn`. Assign it to the default virtual router.
- Enable User Identification on the corp-vpn zone.

On the firewall hosting the external gateway (gpvpn.acme.com):

- Select **Network** > **Interfaces** > **Ethernet** and configure ethernet1/5 as a Layer 3 Ethernet interface with IP address **192.0.2.4** and assign it to the l3-untrust security zone and the default virtual router.
- Create a DNS "A" record that maps IP address 192.0.2.4 to gpvpn.acme.com.
- Select **Network** > **Interfaces** > **Tunnel** and add the tunnel.3 interface and add it to a new zone called `corp-vpn`. Assign it to the default virtual router.
- Enable User Identification on the corp-vpn zone.

On the firewall hosting the internal gateways (california.acme.com and newyork.acme.com):

- Select **Network** > **Interfaces** > **Ethernet** and configure Layer 3 Ethernet interface with IP addresses on the internal network and assign them to the l3-trust security zone and the default virtual router.
- Create a DNS "A" record that maps the internal IP addresses california.acme.com and newyork.acme.com.
- Enable User Identification on the l3-trust zone.

STEP 2 | Purchase and install a GlobalProtect subscription for each firewall hosting a gateway (internal and external) if you have users who will be using the GlobalProtect app on their mobile devices or if you plan to use HIP-enabled security policy.



After you purchase the GlobalProtect subscriptions and receive your activation code, install the GlobalProtect subscriptions on the firewalls hosting your gateways as follows:

1. Select **Device** > **Licenses**.
2. Select **Activate feature using authorization code**.
3. When prompted, enter the **Authorization Code** and then click **OK**.
4. Verify that the license and subscriptions were successfully activated.

Contact your Palo Alto Networks Sales Engineer or Reseller if you do not have the required licenses. For more information on licensing, see About GlobalProtect Licenses.

STEP 3 | Obtain server certificates for the GlobalProtect portal and each GlobalProtect gateway.

In order to connect to the portal for the first time, the end clients must trust the root CA certificate used to issue the portal server certificate.

© 2017 Palo Alto Networks, Inc.

You can use self-signed certificates on the gateways and deploy the root CA certificate to the agents in the client configuration. The best practice is to generate all of the certificates on firewall hosting the portal and deploy them to the gateways.

The recommended workflow is as follows:

1. On the firewall hosting the portal:
    1. Import a server certificate from a well-known, third-party CA.
    2. Create the root CA certificate for issuing self-signed certificates for the GlobalProtect components.
    3. Use the root CA on the portal to generate a self-signed server certificate. Repeat this step for each gateway.
2. On each firewall hosting an internal gateway:
    - Deploy the self-signed server certificates.

STEP 4 | Define how you will authenticate users to the portal and the gateways.

You can use any combination of certificate profiles and/or authentication profiles as necessary to ensure the security for your portal and gateways. Portals and individual gateways can also use different authentication schemes. See the following sections for step-by-step instructions:

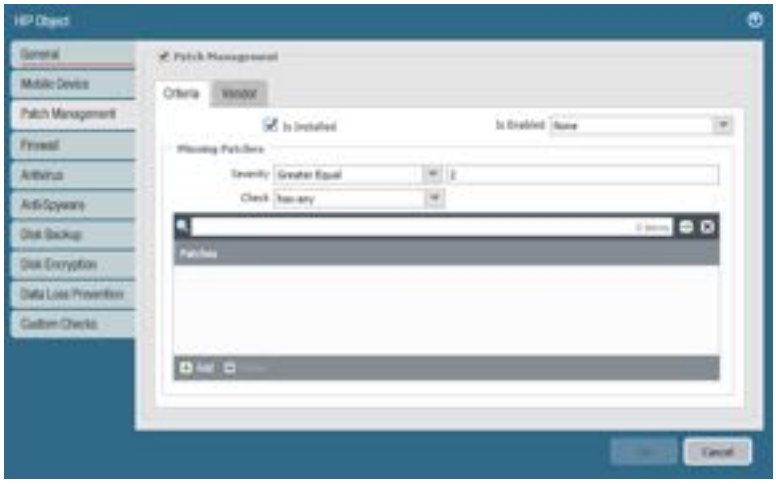
- Set Up External Authentication (authentication profile)
- Set Up Client Certificate Authentication (certificate profile)
- Set Up Two-Factor Authentication (token- or OTP-based)

You will then need to reference the certificate profile and/or authentication profiles you defined in the portal and gateway configurations you define.

STEP 5 | Create the HIP profiles you will need to enforce security policy on gateway access.

See Host Information for more information on HIP matching.

1. Create the HIP objects to filter the raw host data collected by the agents. For example, if you are interested in preventing users that are not up to date with required patches, you might create a HIP object to match on whether the patch management software is installed and that all patches with a given severity are up to date.



2. Create the HIP profiles that you plan to use in your policies.

    For example, if you want to ensure that only Windows users with up-to-date patches can access your internal applications, you might attach the following HIP profile that will match hosts that do NOT have a missing patch:

© 2017 Palo Alto Networks, Inc.



STEP 6 | Configure the internal gateways.

Select **Network** > **GlobalProtect** > **Gateways** and add the following settings:

- **Interface**
- **IP Address**
- **Server Certificate**
- **Authentication Profile** and/or **Configuration Profile**

Notice that it is not necessary to configure the client configuration settings in the gateway configurations (unless you want to set up HIP notifications) because tunnel connections are not required. See Configure a GlobalProtect Gateway for step-by-step instructions on creating the gateway configurations.

STEP 7 | Configure the GlobalProtect Portals.

Although this example shows how to create a single client configuration to be deployed to all agents, you could choose to create separate configurations for different uses and then deploy them based on user/group name and/or the operating system the agent/app is running on (Android, iOS, Mac, or Windows).

Select **Network** > **GlobalProtect** > **Portals** and add the following configuration:

1. Set Up Access to the GlobalProtect Portal:

   Interface—**`ethernet1/2`**

   IP Address—**`10.31.34.13`**

   Server Certificate—**`GP-server-cert.pem issued by GoDaddy`** with **`CN=gp.acme.com`**
2. Define the GlobalProtect Client Authentication Configurations:

   Internal Host Detection—**`enabled`**

   Use single sign-on—**`enabled`**

   Connect Method—**`User-logon (Always On)`**

   External Gateway Address—**`gpvpn.acme.com`**

   Internal Gateway Address—**`california.acme.com, newyork.acme.com`**

   User/User Group—any
3. **Commit** the portal configuration.

STEP 8 | Deploy the GlobalProtect Agent Software.

Select **Device** > **GlobalProtect Client**.

In this example, use the procedure to Host Agent Updates on the Portal.

© 2017 Palo Alto Networks, Inc.

STEP 9 | Create security policy rules on each gateway to safely enable access to applications for your VPN users.

- Create security policy (**Policies** > **Security**) to enable traffic flow between the corp-vpn zone and the l3-trust zone.
- Create HIP-enabled and user/group-based policy rules to enable granular access to your internal datacenter resources.
- For visibility, create rules that allow all of your users web-browsing access to the l3-untrust zone, using the default security profiles to protect you from known threats.



STEP 10 | Save the GlobalProtect configuration.

Click **Commit** on the portal and all gateways.

© 2017 Palo Alto Networks, Inc.

GlobalProtect Architecture

This section outlines an example reference architecture for deploying GlobalProtect™ which secures internet traffic and provides secure access to corporate resources.

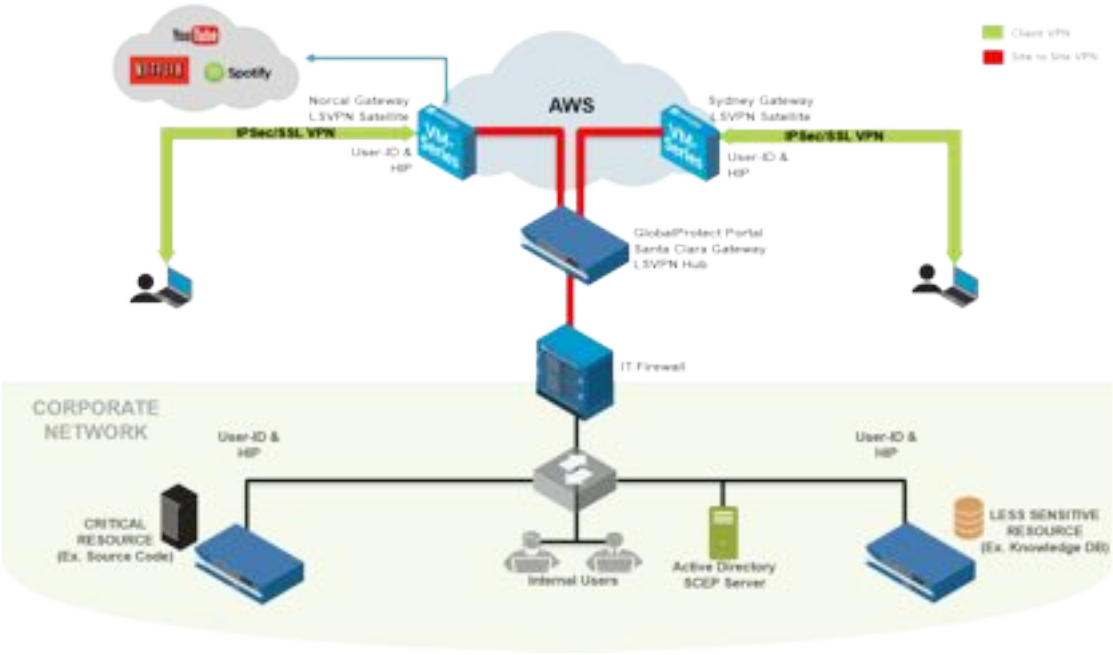
The reference architecture and guidelines described in this section provide a common deployment scenario. Before adopting this architecture, identify your corporate security, infrastructure manageability, and end user experience requirements and deploy GlobalProtect based on those requirements.

Although the requirements may be different for each enterprise, you can leverage the common principles and design considerations outlined in this document along with best practice configuration guidelines to meet your enterprise security needs.

> GlobalProtect Reference Architecture Topology on page 233
> GlobalProtect Reference Architecture Features on page 235
> GlobalProtect Reference Architecture Configurations on page 237

© 2017 Palo Alto Networks, Inc.

GlobalProtect Reference Architecture Topology



- GlobalProtect Portal on page 233
- GlobalProtect Gateways on page 233

GlobalProtect Portal

In this topology, a PA-3020 in the co-location space functions as a GlobalProtect portal.

Employees and contractors can authenticate to the portal using two-factor authentication (2FA) consisting of Active Directory (AD) credentials and a one-time password (OTP). The portal deploys GlobalProtect client configurations based on user and group membership and operating system.

By configuring a separate portal client configuration that applies to a small group or set of pilot users, you can test features before rolling them out to a wider user base. Any client configuration containing new features—such as the Enforce GlobalProtect or Simple Certificate Enrollment Protocol (SCEP) features which were made available with PAN-OS 7.1 and content updates that followed—is enabled in the pilot configuration first and validated by those pilot users, before it is made available to other users.

The GlobalProtect portal also pushes configurations to GlobalProtect satellites. This configuration includes the GlobalProtect gateways to which satellites can connect and establish a site-to-site tunnel.

GlobalProtect Gateways

The PA-3020 in the co-location space (mentioned previously) also doubles as a GlobalProtect gateway (the Santa Clara Gateway). 10 additional gateways are deployed in Amazon Web Services (AWS) and the Microsoft Azure public cloud. The regions or POP locations where these AWS and Azure gateways are deployed are based on the distribution of employees across the globe.

- **Santa Clara Gateway**—Employees and contractors can authenticate to the Santa Clara Gateway (PA-3020 in the co-location space) using 2FA. This gateway requires users to provide their Active

© 2017 Palo Alto Networks, Inc.

Directory credentials and their OTP. Because this gateway protects sensitive resources, it is configured as a manual-only gateway. As a result, users do not connect to this gateway automatically and must manually choose to connect to this gateway. For example, when users connect to AWS-Norcal, which is not a manual-only gateway, some sensitive internal resources are not accessible. The user must then manually switch to and authenticate with the Santa Clara Gateway to access these resources.

In addition, the Santa Clara Gateway is configured as a Large Scale VPN (LSVPN) tunnel termination point for all satellite connections from gateways in AWS and Azure. The Santa Clara Gateway is also configured to set up an Internet Protocol Security (IPSec) tunnel to the IT firewall in corporate headquarters. This is the tunnel that provides access to resources in the corporate headquarters.

- **Gateways in Amazon Web Services and Microsoft Azure**—This gateway requires 2FA: a client certificate and Active Directory credentials. The GlobalProtect portal distributes the client certificate that is required to authenticate with these gateways using the GlobalProtect SCEP feature.

  These gateways in the public cloud also act as GlobalProtect satellites. They communicate with the GlobalProtect portal, download the satellite configuration, and establish a site-to-site tunnel with the Santa Clara Gateway. GlobalProtect satellites initially authenticate using serial number, and subsequently authenticate using certificates.
- **Gateways Inside Corporate Headquarters**—Within the corporate headquarters, three firewalls function as GlobalProtect gateways. These are internal gateways and do not require endpoints to set up a tunnel. Users authenticate to these gateways using their Active Directory credentials. These internal gateways use GlobalProtect to identify the User-ID and to collect Host Information Profile (HIP) from the endpoints.

  To make the end user experience as seamless as possible, you can configure these internal gateways to authenticate users using certificates provisioned by SCEP or using Kerberos service tickets.

© 2017 Palo Alto Networks, Inc.

GlobalProtect Reference Architecture Features

- End User Experience on page 235
- Management and Logging on page 235
- Monitoring and High Availability on page 236

End User Experience

End users who are remote (not inside the corporate network) connect to one of the gateways in AWS or Azure. When you configure the GlobalProtect portal client configuration, assign equal priority to the gateways. With this configuration, the gateway to which users connect depends on the SSL response time of each gateway measured on the endpoint during the tunnel set-up time.

For example, a user in Australia would typically connect to the AWS-Sydney Gateway. Once the user is connected to AWS-Sydney, GlobalProtect client tunnels all traffic from the endpoint to the AWS-Sydney firewall for inspection. GlobalProtect sends traffic to public internet sites directly via the AWS-Sydney Gateway and tunnels traffic to corporate resources through a site-to-site tunnel between the AWS-Sydney Gateway and the Santa Clara Gateway, and then through an IPSec site-to-site tunnel to the corporate headquarters. This architecture is designed to reduce any latency the user may experience when accessing the internet. If the AWS-Sydney Gateway (or any gateway closer to Sydney) was unreachable, the GlobalProtect client would back-haul the internet traffic to the firewall in the corporate headquarters and cause latency issues.

Active directory servers reside inside the corporate network. When remote end users authenticate, the GlobalProtect client sends authentication requests through the site-to-site tunnel in AWS/Azure to the Santa Clara Gateway. The gateway then forwards the request through an IPSec site-to-site tunnel to the Active Directory Server in corporate headquarters.



To reduce the time it takes for remote user authentication and tunnel set up, consider replicating the Active Directory Server and making it available in AWS.

End users inside the corporate network authenticate to the three internal gateways immediately after they log in; The GlobalProtect client sends the HIP report to these internal gateways. When users are inside the office on the corporate network, they must meet the User-ID and HIP requirements to access any resource at work.

Management and Logging

In this deployment, you can manage and configure all firewalls from Panorama, which is deployed in the co-location space.

To provide consistent security, all firewalls in AWS and Azure use the same security policies and configurations. To simplify configuration of the gateways, Panorama also uses one device group and one template. In this deployment, all gateways forward all logs to Panorama. This enables you to monitor network traffic or troubleshoot issues from a central location instead of requiring you to log in to each firewall.

When software updates are required, you can use Panorama to deploy the software updates to all firewalls. Panorama first upgrades one or two firewalls and verifies whether the upgrade was successful before updating the remaining firewalls.

© 2017 Palo Alto Networks, Inc.

Monitoring and High Availability

To monitor the firewalls in this deployment, you can use Nagios, an open-source server, network, and log monitoring software. Configure Nagios to periodically verify the response from the portal and the gateways' pre-login page and send an alert if the response does not match the expectations. You can also configure GlobalProtect Simple Network Management Protocol (SNMP) Management Information Base (MIB) objects to monitor gateway usage.

In this deployment there is only one instance of the GlobalProtect portal. If the portal becomes unavailable, new users (who have never connected to the portal before) will not be able to connect to GlobalProtect. However, existing users can use the cached portal client configuration to connect to one of the gateways.

Multiple virtual machine (VM) firewalls in AWS configured as GlobalProtect gateways provide gateway redundancy. Therefore, configuring gateways as a high availability (HA) pair is not required.

© 2017 Palo Alto Networks, Inc.

GlobalProtect Reference Architecture Configurations

To align your deployment with the reference architecture, review the following configuration checklists.

- Gateway Configuration on page 237
- Portal Configuration on page 237
- Policy Configurations on page 237

Gateway Configuration

- Disable split tunneling. To do this, ensure there are no Access Routes specified in **Agent** > **Client Settings** > **Split Tunnel** settings. See Configure a GlobalProtect Gateway on page 72.
- Enable **No direct access to local network** in **Agent** > **Client Settings** > **Split Tunnel**. See Configure a GlobalProtect Gateway on page 72.
- Enable the gateway to **Accept cookie for authentication override**. See Configure a GlobalProtect Gateway on page 72.

Portal Configuration

- Configure the **Connect Method** as **Always-on (User logon)**. See Customize the GlobalProtect Agent.
- Set **Use Single Sign-On** (Windows only) to **Yes**. See Customize the GlobalProtect Agent.
- Configure the portal to **Save User Credentials** (set the value to **Yes**). See Define the GlobalProtect Agent Configurations.
- Enable the portal to **Accept cookie for authentication override**. See Define the GlobalProtect Agent Configurations.
- Configure the **Cookie Lifetime** as 20 hours. See Define the GlobalProtect Agent Configurations.
- **Enforce GlobalProtect** for network access. See Customize the GlobalProtect Agent.
- When **Enforce GlobalProtect for Network Access** is enabled, allow users to disable the GlobalProtect agent with a passcode. See Customize the GlobalProtect Agent.
- Configure **Internal Host Detection**. See Define the GlobalProtect Agent Configurations.
- Enable the **Collect HIP Data** option in Data Collection. See Define the GlobalProtect Agent Configurations.
- Distribute and install the SSL Forward Proxy CA certificate used for SSL Decryption. See Define the GlobalProtect Agent Configurations.

Policy Configurations

- Configure all firewalls to use security policies and profiles based on the Best Practice Internet Gateway Security Policy. In this reference deployment, this includes the Santa Clara Gateway in the co-location space and gateways in the AWS/Azure public cloud.
- Enable SSL Decryption on all gateways in AWS and Azure.
- Configure Policy-Based Forwarding rules for all gateways in AWS to forward traffic to certain websites through the Santa Clara Gateway. This ensures that sites like www.stubhub.com and www.lowes.com that block traffic from AWS IP address ranges are still accessible when users connect to gateways in AWS.

© 2017 Palo Alto Networks, Inc.

GlobalProtect Cryptography

- About GlobalProtect Cipher Selection
- Cipher Exchange Between the GlobalProtect Agent and Gateway
- GlobalProtect Cryptography References
- Ciphers Used to Set Up IPSec Tunnels
- SSL APIs

© 2017 Palo Alto Networks, Inc.

About GlobalProtect Cipher Selection

GlobalProtect supports both IPSec and SSL tunnel modes. GlobalProtect also supports the ability to enable and require the GlobalProtect agent to always attempt to set up IPSec tunnel first before falling back to SSL tunnel. With an IPSec tunnel, the GlobalProtect agent uses SSL/TLS to exchange encryption and authentication algorithms and the keys. The selection of cipher suite that GlobalProtect uses to secure the SSL/TLS tunnel depend on:

- **SSL/TLS versions accepted by the gateway**—The GlobalProtect portal and gateways can restrict the list of cipher suites available for the client using SSL/TLS profiles. On the firewall, you create the SSL/TLS profile by specifying the certificate and the allowed protocol versions and associate that to the GlobalProtect portal and gateway.
- **Algorithm of the server certificate of the gateway**—The operating system of the endpoint determines what cipher suites the GlobalProtect agent includes in its Client Hello message. As long as the GlobalProtect agent includes the cipher suite that gateway prefers to use, the gateway will select that cipher suite for the SSL session. The order of cipher suites within the Client Hello message does not affect the cipher suite selection: The gateway selects the cipher suite based on the SSL/TLS service profile and the algorithm of the gateway server certificate and its preferred list. You select the service profile from the GlobalProtect gateway authentication configuration.

© 2017 Palo Alto Networks, Inc.

Cipher Exchange Between the GlobalProtect Agent and Gateway

The following figure displays the exchange of ciphers between GlobalProtect gateways and GlobalProtect agents when creating the VPN tunnel.

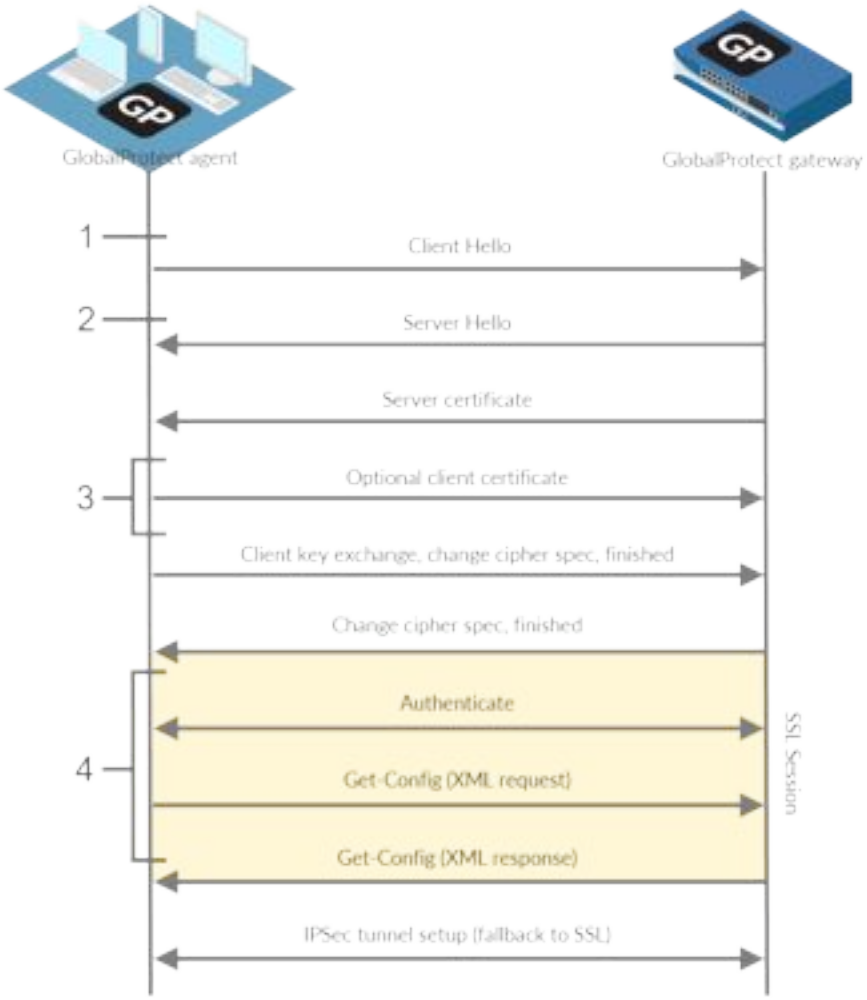


Figure 8: Cipher Exchange Between the Agent and the Gateway

The following table describes these stages in more detail.

Table 7: Cipher Exchange Between the Agent and Gateway

Communication Stage	Description
1. Client Hello	The agent proposes a list of cipher suites depending on the OS of the endpoint.
2. Server Hello	The gateway selects the cipher suite proposed by the agent. When selecting the ciphers to set up the tunnel, the gateway ignores both the number and

© 2017 Palo Alto Networks, Inc.

Communication Stage	Description
	order of cipher suites proposed by the agent and instead relies on the SSL/TLS versions and algorithm of the gateway server certificate and its preferred list (as described in About GlobalProtect Cipher Selection).
3. Optional Client Certificate	The gateway can optionally request a client certificate from the agent to use to trust the identity of the user or endpoint.
4. SSL Session	After setting up the SSL/TLS session, the agent authenticates with the gateway and requests the gateway configuration (Get-Config-Request). To request the configuration, the agent proposes the encryption and authentication algorithms and other settings such as preferred IP address for the tunnel interface. The gateway responds to the request and selects the encryption and authentication algorithm to use based on the configuration of the GlobalProtect IPSec Crypto Profile (Get-Config-Response).

The following table displays an example of the cipher exchange between an agent on a Mac endpoint and the gateway.

Table 8: Example: Cipher Exchange for Mac Endpoints

Communication Stage	Example: Mac Endpoints
1. Client Hello	TLS 1.2 37 Cipher Suites (Reference: TLS Ciphers Supported by GlobalProtect Agents on Mac Endpoints)
2. Server Hello	• When GlobalProtect uses an ECDSA certificate and TLS 1.2 is accepted, the SSL session uses ECDSA-AES256-CBC-SHA. • When GlobalProtect uses an RSA certificate and TLS 1.2 is accepted, the SSL session uses RSA-AES256-CBC-SHA256.
3. Optional Client Certificate	Client certificates signed with ECDSA or RSA and using SHA1, SHA256, or SHA384

4. SSL Session	• SSL Session uses ECDSA-AES256-CBC-SHA or RSA-AES256-CBC-SHA256
	• Get-Config-Request
	  • Encryption—AES-256-GCM, AES-128-GCM, AES-128-CBC
	  • Authentication—SHA1 and OS type, Preferred IP address etc
	• Get-Config-Response
	  • Client to server, and server to client SPIs, encryption keys, and authentication keys
	  • Tunnel type, ports, split tunnel mode, IP, and DNS etc

© 2017 Palo Alto Networks, Inc.

GlobalProtect Cryptography References

- Reference: GlobalProtect Agent Cryptographic Functions
- TLS Cipher Suites Supported by GlobalProtect Agents
- TLS Cipher Suites Supported by GlobalProtect Gateways in PAN-OS 8.0

Reference: GlobalProtect Agent Cryptographic Functions

The GlobalProtect agent uses the OpenSSL library 1.0.1h to establish secure communication with the GlobalProtect portal and GlobalProtect gateways. The following table lists each GlobalProtect agent function that requires a cryptographic function and the cryptographic keys the GlobalProtect agent uses:

Crypto Function	Key	Usage
Winhttp (Windows) and NSURLConnection (MAC) aes256-sha	Dynamic key negotiated between the GlobalProtect agent and the GlobalProtect portal and/or gateway for establishing the HTTPS connection.	Used to establish the HTTPS connection between the GlobalProtect agent and the GlobalProtect portal and GlobalProtect gateway for authentication.
OpenSSL aes256-sha	Dynamic key negotiated between the GlobalProtect agent and the GlobalProtect gateway during the SSL handshake.	Used to establish the SSL connection between the GlobalProtect agent and the GlobalProtect gateway for HIP report submission, SSL tunnel negotiation, and network discovery.
IPSec encryption and authentication aes-128-sha1, aes-128-cbc, aes-128-gcm, and aes-256-gcm	The session key sent from the GlobalProtect gateway.	Used to establish the IPSec tunnel between the GlobalProtect agent and the GlobalProtect gateway. Use the strongest algorithm supported by your network (AES-GCM is recommended). To provide data integrity and authenticity protection, the aes-128-cbc cipher requires the sha1 authentication algorithm. Because AES-GCM encryption algorithms (aes-128-gcm and aes-256-gcm) natively provide ESP integrity protection, the sha1 authentication algorithm is ignored for these ciphers even though it is required during configuration.

© 2017 Palo Alto Networks, Inc.

TLS Cipher Suites Supported by GlobalProtect Agents

The following tables provide examples of TLS ciphers supported on GlobalProtect agents installed on various endpoint operating systems. The lists are not exhaustive for all supported OSs.

- Reference: TLS Ciphers Supported by GlobalProtect Agents on Mac Endpoints
- Reference: TLS Ciphers Supported by GlobalProtect Agents on Windows 7 Endpoints
- Reference: TLS Ciphers Supported by GlobalProtect Agents on Android 6.0.1 Endpoints
- Reference: TLS Ciphers Supported by GlobalProtect Agents on iOS 10.2.1 Endpoints
- Reference: TLS Ciphers Supported by GlobalProtect Agents on Chromebooks

Reference: TLS Ciphers Supported by GlobalProtect Agents on Mac Endpoints

TLS Ciphers Supported by GlobalProtect Agents on Mac Endpoints	
TLS_EMPTY_RENEGOTIATION_INFO_SCSV (0x00ff)	TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384 (0xc02a)
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384 (0xc024)	TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256 (0xc029)
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256 (0xc023)	TLS_ECDH_RSA_WITH_AES_256_CBC_SHA (0xc00f)
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA (0xc00a)	TLS_ECDH_RSA_WITH_AES_128_CBC_SHA (0xc00e)
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA (0xc009)	TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA (0xc00d)
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA (0xc008)	TLS_DHE_RSA_WITH_AES_256_CBC_SHA256 (0x006b)
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (0xc028)	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256 (0x0067)
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (0xc027)	TLS_DHE_RSA_WITH_AES_256_CBC_SHA (0x0039)
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (0xc014)	TLS_DHE_RSA_WITH_AES_128_CBC_SHA (0x0033)
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (0xc013)	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA (0x0016)
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA (0xc012)	TLS_RSA_WITH_AES_256_CBC_SHA256 (0x003d)
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384 (0xc026)	TLS_RSA_WITH_AES_128_CBC_SHA256 (0x003c)
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256 (0xc025)	TLS_RSA_WITH_AES_256_CBC_SHA (0x0035)
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA (0xc005)	TLS_RSA_WITH_AES_128_CBC_SHA (0x002f)
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA (0xc004)	TLS_RSA_WITH_3DES_EDE_CBC_SHA (0x000a)
	TLS_ECDHE_ECDSA_WITH_RC4_128_SHA (0xc007)
	TLS_ECDHE_RSA_WITH_RC4_128_SHA (0xc011)

© 2017 Palo Alto Networks, Inc.

TLS Ciphers Supported by GlobalProtect Agents on Mac Endpoints	
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA (0xc003)	TLS_ECDH_ECDSA_WITH_RC4_128_SHA (0xc002)
	TLS_ECDH_RSA_WITH_RC4_128_SHA (0xc00c)
	TLS_RSA_WITH_RC4_128_SHA (0x0005)
	TLS_RSA_WITH_RC4_128_MD5 (0x0004)

Reference: TLS Ciphers Supported by GlobalProtect Agents on Windows 7 Endpoints

TLS Ciphers Supported by GlobalProtect Agents on Windows 7 Endpoints	
TLS_EMPTY_RENEGOTIATION_INFO_SCSV (0x00ff)	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (0xc02f)
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_ SHA384 (0xc02c)	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (0xc028)
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_ SHA256 (0xc02b)	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (0xc027)
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_ SHA384 (0xc024)	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (0xc014)
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_ SHA256 (0xc023)	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (0xc013)
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA (0xc00a)	TLS_RSA_WITH_AES_256_GCM_SHA384 (0x009d)
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA (0xc009)	TLS_RSA_WITH_AES_128_GCM_SHA256 (0x009c)
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (0xc030)	TLS_RSA_WITH_AES_256_CBC_SHA256 (0x003d)
	TLS_RSA_WITH_AES_128_CBC_SHA256 (0x003c)
	TLS_RSA_WITH_AES_256_CBC_SHA (0x0035)
	TLS_RSA_WITH_AES_128_CBC_SHA (0x002f)

Reference: TLS Ciphers Supported by GlobalProtect Agents on Android 6.0.1 Endpoints

The GlobalProtect agent for Android 6.0.1 supports 20 cipher suites.

TLS Ciphers Supported by GlobalProtect Agents on Android 6.0.1 Endpoints	
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_ SHA256 (0xc02b)	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (0xc014)
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_ SHA384 (0xc02c)	TLS_DHE_RSA_WITH_AES_128_CBC_SHA (0x0033)

© 2017 Palo Alto Networks, Inc.

TLS Ciphers Supported by GlobalProtect Agents on Android 6.0.1 Endpoints	
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (0xc02f)	TLS_DHE_RSA_WITH_AES_256_CBC_SHA (0x0039)
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (0xc030)	TLS_ECDHE_ECDSA_WITH_RC4_128_SHA (0xc007)
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 (0x009e)	TLS_ECDHE_RSA_WITH_RC4_128_SHA (0xc011)
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 (0x009f)	TLS_RSA_WITH_AES_128_GCM_SHA256 (0x009c)
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA (0xc009)	TLS_RSA_WITH_AES_256_GCM_SHA384 (0x009d)
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA (0xc00a)	TLS_RSA_WITH_AES_128_CBC_SHA (0x002f)
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (0xc013)	TLS_RSA_WITH_AES_256_CBC_SHA (0x0035)
	TLS_RSA_WITH_RC4_128_SHA (0x0005)
	TLS_EMPTY_RENEGOTIATION_INFO_SCSV (0x00ff)

Reference: TLS Ciphers Supported by GlobalProtect Agents on iOS 10.2.1 Endpoints

The GlobalProtect app for iOS 10.2.1 supports 19 cipher suites.

TLS Ciphers Supported by GlobalProtect Agents on iOS 10.2.1 Endpoints	
TLS_EMPTY_RENEGOTIATION_INFO_SCSV (0x00ff)	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (0xc028)
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_ SHA384 (0xc02c)	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (0xc027)
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_ SHA256 (0xc02b)	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (0xc014)
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_ SHA384 (0xc024)	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (0xc013)
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_ SHA256 (0xc023)	TLS_RSA_WITH_AES_256_GCM_SHA384 (0x009d)
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA (0xc00a)	TLS_RSA_WITH_AES_128_GCM_SHA256 (0x009c)
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA (0xc009)	TLS_RSA_WITH_AES_256_CBC_SHA256 (0x003d)
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (0xc030)	TLS_RSA_WITH_AES_128_CBC_SHA256 (0x003c)
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (0xc02f)	TLS_RSA_WITH_AES_256_CBC_SHA (0x0035)
	TLS_RSA_WITH_AES_128_CBC_SHA (0x002f)

© 2017 Palo Alto Networks, Inc.

Reference: TLS Ciphers Supported by GlobalProtect Agents on Chromebooks

The GlobalProtect app for Chrome OS 55.0.2883 supports 91 cipher suites.

TLS Ciphers Supported by GlobalProtect Agents on Chromebooks (Chrome OS 55.0.2883)	
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (0xc030)	TLS_DH_DSS_WITH_CAMELLIA_256_CBC_SHA (0x0085)
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_ SHA384 (0xc02c)	TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384 (0xc032)
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (0xc028)	TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384 (0xc02e)
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_ SHA384 (0xc024)	TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384 (0xc02a)
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (0xc014)	TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384 (0xc026)
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA (0xc00a)	TLS_ECDH_RSA_WITH_AES_256_CBC_SHA (0xc00f)
TLS_DH_DSS_WITH_AES_256_GCM_SHA384 (0x00a5)	TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA (0xc005)
TLS_DHE_DSS_WITH_AES_256_GCM_SHA384 (0x00a3)	TLS_RSA_WITH_AES_256_GCM_SHA384 (0x009d)
TLS_DH_RSA_WITH_AES_256_GCM_SHA384 (0x00a1)	TLS_RSA_WITH_AES_256_CBC_SHA256 (0x003d)
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 (0x009f)	TLS_RSA_WITH_AES_256_CBC_SHA (0x0035)
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256 (0x006b)	TLS_RSA_WITH_CAMELLIA_256_CBC_SHA (0x0084)
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256 (0x006a)	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (0xc02f)
TLS_DH_RSA_WITH_AES_256_CBC_SHA256 (0x0069)	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 (0xc02b)
TLS_DH_DSS_WITH_AES_256_CBC_SHA256 (0x0068)	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (0xc027)
TLS_DHE_RSA_WITH_AES_256_CBC_SHA (0x0039)	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256 (0xc023)
TLS_DHE_DSS_WITH_AES_256_CBC_SHA (0x0038)	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (0xc013)
TLS_DH_RSA_WITH_AES_256_CBC_SHA (0x0037)	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA (0xc009)
TLS_DH_DSS_WITH_AES_256_CBC_SHA (0x0036)	TLS_DH_DSS_WITH_AES_128_GCM_SHA256 (0x00a4)
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA (0x0088)	TLS_DHE_DSS_WITH_AES_128_GCM_SHA256 (0x00a2)

© 2017 Palo Alto Networks, Inc.

TLS Ciphers Supported by GlobalProtect Agents on Chromebooks (Chrome OS 55.0.2883)	
TLS_DHE_DSS_WITH_CAMELLIA_256_CBC_SHA (0x0087) TLS_DH_RSA_WITH_CAMELLIA_256_CBC_SHA (0x0086)	TLS_DH_RSA_WITH_AES_128_GCM_SHA256 (0x00a0) TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 (0x009e)
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256 (0x0067) TLS_DHE_DSS_WITH_AES_128_CBC_SHA256 (0x0040) TLS_DH_RSA_WITH_AES_128_CBC_SHA256 (0x003f) TLS_DH_DSS_WITH_AES_128_CBC_SHA256 (0x003e) TLS_DHE_RSA_WITH_AES_128_CBC_SHA (0x0033) TLS_DHE_DSS_WITH_AES_128_CBC_SHA (0x0032) TLS_DH_RSA_WITH_AES_128_CBC_SHA (0x0031) TLS_DH_DSS_WITH_AES_128_CBC_SHA (0x0030) TLS_DHE_RSA_WITH_SEED_CBC_SHA (0x009a) TLS_DHE_DSS_WITH_SEED_CBC_SHA (0x0099) TLS_DH_RSA_WITH_SEED_CBC_SHA (0x0098) TLS_DH_DSS_WITH_SEED_CBC_SHA (0x0097) TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA (0x0045) TLS_DHE_DSS_WITH_CAMELLIA_128_CBC_SHA (0x0044) TLS_DH_RSA_WITH_CAMELLIA_128_CBC_SHA (0x0043) TLS_DH_DSS_WITH_CAMELLIA_128_CBC_SHA (0x0042) TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256 (0xc031) TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256 (0xc02d) TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256 (0xc029) TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256 (0xc025)	TLS_RSA_WITH_AES_128_CBC_SHA256 (0x003c) TLS_RSA_WITH_AES_128_CBC_SHA (0x002f) TLS_RSA_WITH_SEED_CBC_SHA (0x0096) TLS_RSA_WITH_CAMELLIA_128_CBC_SHA (0x0041) TLS_RSA_WITH_IDEA_CBC_SHA (0x0007) TLS_ECDHE_RSA_WITH_RC4_128_SHA (0xc011) TLS_ECDHE_ECDSA_WITH_RC4_128_SHA (0xc007) TLS_ECDH_RSA_WITH_RC4_128_SHA (0xc00c) TLS_ECDH_ECDSA_WITH_RC4_128_SHA (0xc002) TLS_RSA_WITH_RC4_128_SHA (0x0005) TLS_RSA_WITH_RC4_128_MD5 (0x0004) TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA (0xc012) TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA (0xc008) TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA (0x0016) TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA (0x0013) TLS_DH_RSA_WITH_3DES_EDE_CBC_SHA (0x0010) TLS_DH_DSS_WITH_3DES_EDE_CBC_SHA (0x000d) TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA (0xc00d) TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA (0xc003) TLS_RSA_WITH_3DES_EDE_CBC_SHA (0x000a) TLS_DHE_RSA_WITH_DES_CBC_SHA (0x0015) TLS_DHE_DSS_WITH_DES_CBC_SHA (0x0012) TLS_DH_RSA_WITH_DES_CBC_SHA (0x000f) TLS_DH_DSS_WITH_DES_CBC_SHA (0x000c)

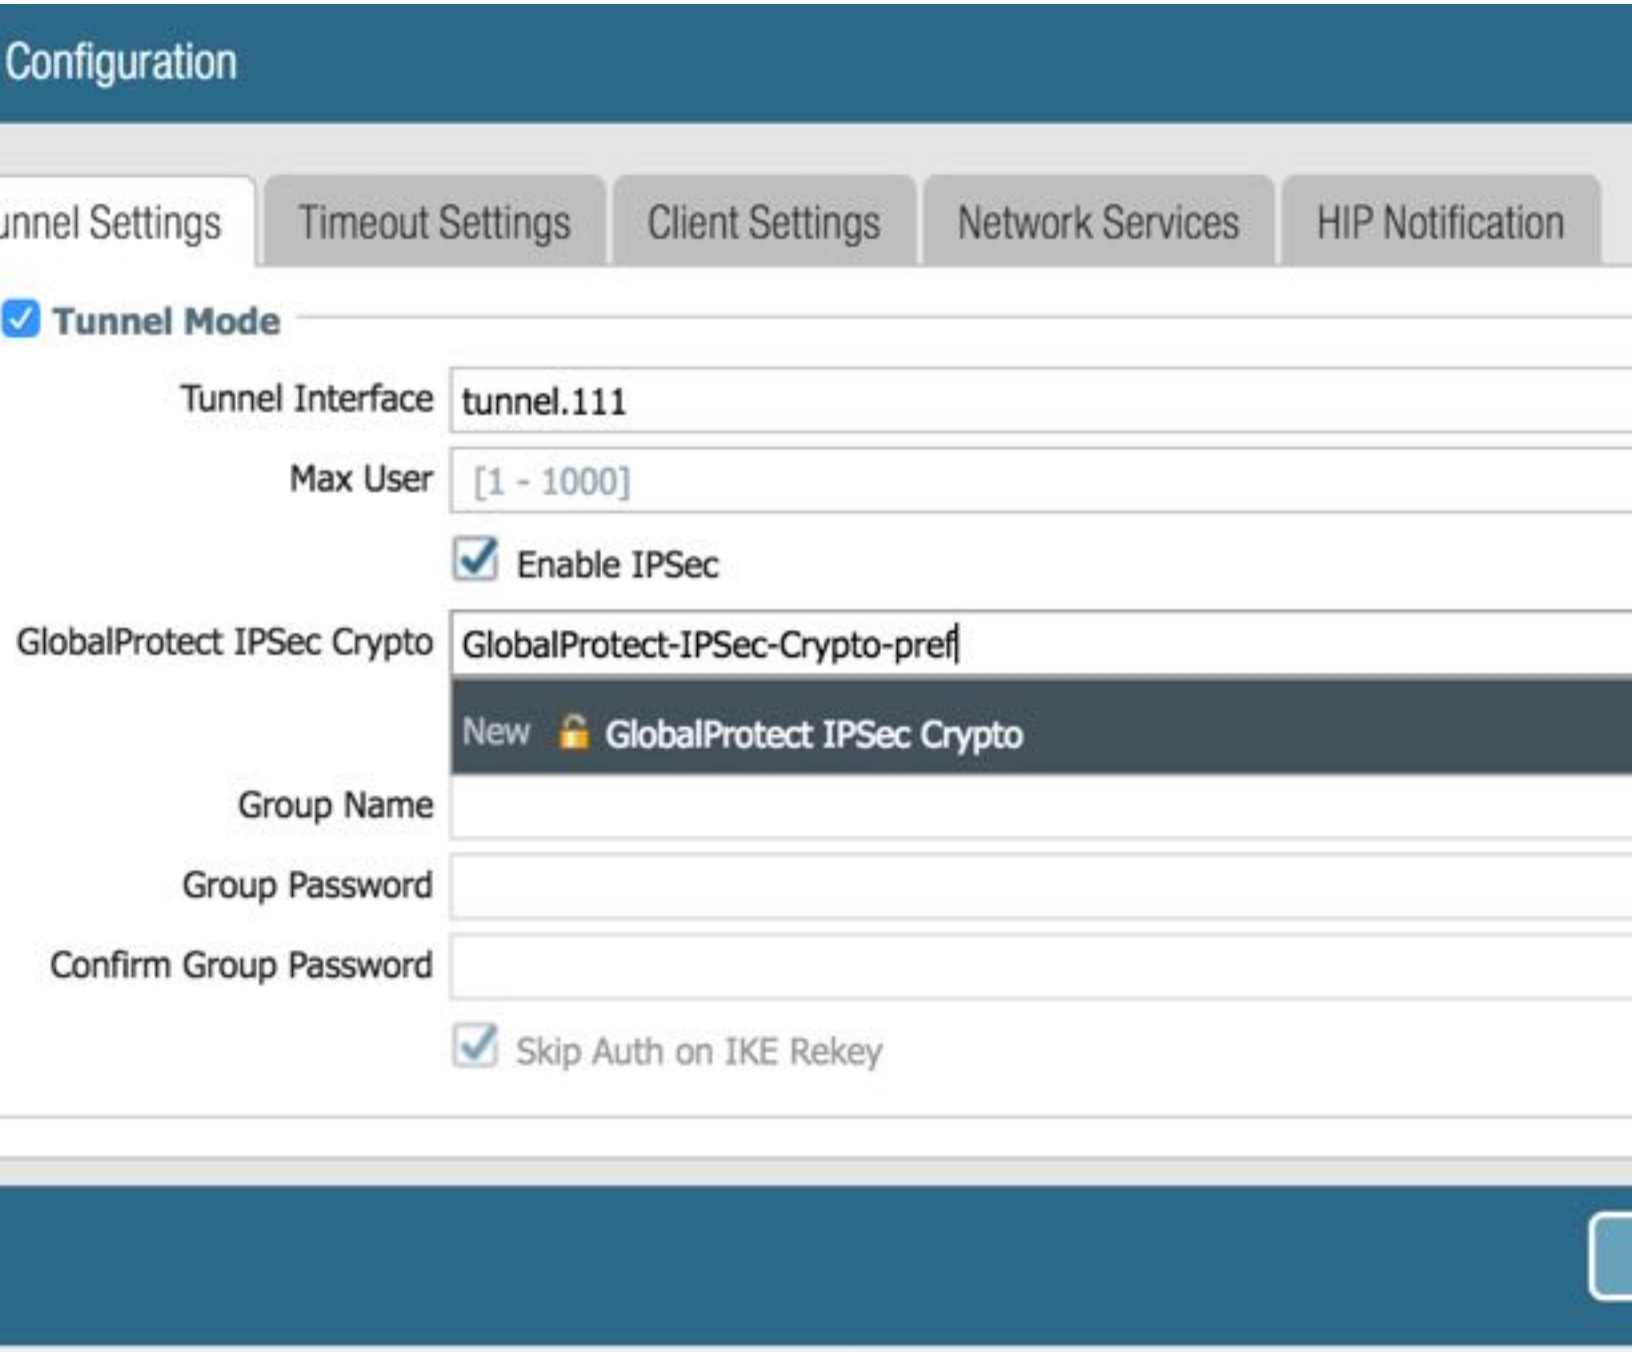
© 2017 Palo Alto Networks, Inc.

TLS Ciphers Supported by GlobalProtect Agents on Chromebooks (Chrome OS 55.0.2883)	
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA (0xc00e) TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA (0xc004) TLS_RSA_WITH_AES_128_GCM_SHA256 (0x009c)	TLS_RSA_WITH_DES_CBC_SHA (0x0009) TLS_EMPTY_RENEGOTIATION_INFO_SCSV (0x00ff)

© 2017 Palo Alto Networks, Inc.

Ciphers Used to Set Up IPSec Tunnels

GlobalProtect can restrict and/or set preferential order for what encryption and authentication algorithm the GlobalProtect agent can use for the IPSec tunnel. The algorithms and preferences are defined in the GlobalProtect IPSec Crypto Profile that you configure when you set up the tunnel settings of the GlobalProtect gateway.



When the GlobalProtect agent sets up an SSL session with GlobalProtect gateway. The cipher suite used for this SSL session is governed by the SSL/TLS profile configured on the gateway and the type of algorithm

© 2017 Palo Alto Networks, Inc.

used by the gateway certificate. After the SSL session is established, the GlobalProtect agent initiates a VPN tunnel setup by requesting the configuration over SSL.

Using the same SSL session, the GlobalProtect gateway responds with the encryption and authentication algorithms, keys, and SPIs the agent should use to set up the IPSec tunnel.



AES-GCM is recommended for more secure requirements. To provide data integrity and authenticity protection, the aes-128-cbc cipher requires the SHA1 authentication algorithm. Because AES-GCM encryption algorithms (aes-128-gcm and aes-256-gcm) natively provide ESP integrity protection, the SHA1 authentication algorithm is ignored for these ciphers even though it is required during configuration.

The GlobalProtect IPSec Crypto Profile that you configured on the gateway determines the encryption and authentication algorithm used to set up the IPSec tunnel. The GlobalProtect gateway responds with the first matching encryption algorithm listed in the IPSec Crypto Profile that matches the agent's proposal.

The GlobalProtect agent then attempts to set up a tunnel based on the response from the gateway.

© 2017 Palo Alto Networks, Inc.

SSL APIs

GlobalProtect agent uses both OpenSSL and native system APIs for doing the SSL Handshake. Operations such as the GlobalProtect gateway latency measurement—used by GlobalProtect to select the best gateway—gateway logout, and sending the HIP check message and report are all performed over an SSL session that is set up using OpenSSL library. While operations like gateway pre-login, login and get-config are all done over the SSL session that is set up using the native system API.

© 2017 Palo Alto Networks, Inc.

PA-5200 SERIES

Palo Alto Networks® PA-5200 Series of next-generation firewall appliances comprises the PA-5280, PA-5260, PA-5250 and PA-5220, which are ideal for high-speed data center, internet gateway and service provider deployments. The PA-5200 Series delivers up to 68 Gbps of throughput using dedicated processing and memory for the key functional areas of networking, security, threat prevention and management.

Key Security Features

Classifies all applications, on all ports, all the time

- Identifies the application, regardless of port, encryption (SSL or SSH) or evasive technique employed.
- Uses the application, not the port, as the basis for all of your safe enablement policy decisions: allow, deny, schedule, inspect and apply traffic-shaping.
- Categorizes unidentified applications for policy control, threat forensics or App-ID™ technology development.

Enforces security policies for any user, at any location

- Deploys consistent policies to local and remote users running on the Windows®; Apple® Mac® OS X®, macOS® and iOS; Linux; and Android® platforms.
- Enables agentless integration with Microsoft® Active Directory® and Terminal Services, LDAP, Novell® eDirectory™ and Citrix®.
- Easily integrates your firewall policies with 802.1X wireless, proxies, NAC solutions and any other source of user identity information.

Prevents known and unknown threats

- Blocks a range of known threats, including exploits, malware and spyware, across all ports, regardless of common threat-evasion tactics employed.
- Limits the unauthorized transfer of files and sensitive data, and safely enables non-work-related web surfing.
- Identifies unknown malware, analyzes it based on hundreds of malicious behaviors, and then automatically creates and delivers protection.



PA-5200 Series

The controlling element of the PA-5200 Series is PAN-OS® security operating system, which natively classifies all traffic, inclusive of applications, threats and content, and then ties that traffic to the user, regardless of location or device type. The application, content and user – in other words, the elements that run your business – are then used as the basis of your security policies, resulting in an improved security posture and a reduction in incident response time.

Performance and Capacities	PA-5280	PA-5260	PA-5250	PA-5220
Firewall throughput[1] (App-ID enabled)	68 Gbps	68 Gbps	39 Gbps	18 Gbps
Threat Prevention throughput[2]	30 Gbps	30 Gbps	20 Gbps	9 Gbps
IPsec VPN throughput	24 Gbps	24 Gbps	16 Gbps	8 Gbps
Max sessions	64,000,000	32,000,000	8,000,000	4,000,000
New sessions per second[3]	462,000	462,000	348,000	171,000
Virtual systems (base/max[4])	25/225	25/225	25/125	10/20

1. Firewall throughput measured with App-ID and logging enabled utilizing 64KB HTTP transactions
2. Threat Prevention throughput measured with App-ID, IPS, antivirus, anti-spyware, WildFire and logging enabled utilizing 64KB HTTP transactions
3. New sessions per second measured with application-override utilizing 1 byte HTTP transactions
4. Adding virtual systems over base quantity requires a separately purchased license

Networking Features

Interface Modes

- L2, L3, tap, virtual wire (transparent mode)

Routing

- OSPFv2/v3 with graceful restart, BGP with graceful restart, RIP, Static routing
- Policy-based forwarding
- Point-to-point protocol over Ethernet (PPPoE) and DHCP supported for dynamic address assignment
- Multicast: PIM-SM, PIM-SSM, IGMP v1, v2 and v3
- Bidirectional Forwarding Detection (BFD)

IPv6

- L2, L3, tap, virtual wire (transparent mode)
- Features: App-ID, User-ID™, Content-ID™, WildFire® and SSL decryption
- SLAAC

IPsec VPN

- Key exchange: manual key, IKEv1 and IKEv2 (pre-shared key, certificate-based authentication)
- Encryption: 3DES, AES (128-bit, 192-bit, 256-bit)
- Authentication: MD5, SHA-1, SHA-256, SHA-384, SHA-512
- GlobalProtect™ large-scale VPN (LSVPN) for simplified configuration and management

VLANs

- 802.1q VLAN tags per device/per interface: 4,094/4,094
- Aggregate interfaces (802.3ad), LACP

Network Address Translation

- NAT modes (IPv4): static IP, dynamic IP, dynamic IP and port (port address translation)
- NAT64, NPTv6
- Additional NAT features: dynamic IP reservation, tunable dynamic IP and port oversubscription

High Availability

Modes: active/active, active/passive

Failure detection: path monitoring, interface monitoring

Mobile Network Infrastructure

- GTP Security
- SCTP Security

To view additional information about the features and associated capacities of the PA-5200 Series, please visit www.paloaltonetworks.com/products.

Hardware Specifications

I/O

PA-5280 | PA-5260 | PA-5250 – (4) 100/1000/10G Cu, (16) Gig/10Gig SFP/SFP+, (4) 40G/100G QSFP28

PA-5220 – (4)100/1000/10G Cu, (16) Gig/10Gig SFP/SFP+, (4) 40G QSFP+

Management I/O

PA-5280 | PA-5260 | PA-5250 – (2) 10/100/1000, (1) 40G/100G QSFP28 HA, (1) 10/100/1000 out-of-band management, (1) RJ45 console port

PA-5220 – (2) 10/100/1000, (1) 40G QSFP+ HA, (1) 10/100/1000 out-of-band management, (1) RJ45 console port

Storage Capacity

240GB SSD, RAID1, system storage

2TB HDD, RAID1, log storage

Power Supply (Avg/Max Power Consumption)

571/685 Watts

Max BTU/hr

2,340

Power Supplies (Base/Max)

1:1 fully redundant (2/2)

AC Input Voltage (input Hz)

100-240VAC (50-60Hz)

AC Power Supply Output

1,200 Watts/power supply

Max Current

AC power supplies — 8.5A@100VAC, 3.6A@240VAC
DC power supplies — 19A@-40VDC, 12.7A@-60VDC

Max Inrush Current

AC power supplies — 50A@230VAC, 50A@120VAC
DC power supplies — 200A@72VDC

Mean Time Between Failure (MTBF)

9.23 years

Rack Mount (Dimensions)

3U, 19" standard rack
5.25" H x 20.5" D x 17.25" W (13.33cm x 52.07cm x 43.81cm)

Weight (Stand-Alone Device/As Shipped)

46 lbs (20.87 kg)/62 lbs (28.13 kg)

Safety

cCSAus, CB IEC 60950-1

EMI

FCC Class A, CE Class A, VCCI Class A

Certifications

See https://www.paloaltonetworks.com/company/certifications.html

Environment

Operating Temperature: 32° to 122° F, 0° to 50° C
Non-Operating Temperature: -4° to 158° F, -20° to 70° C

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

© 2018 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.
pa-5200-series-ds-041718

Palo Alto Networks Platform Specifications and Features Summary

Performance and Capacities[1]	PA-7080 System[2]	PA-7050 System[2]	PA-5280	PA-5260	PA-5250	PA-5220
Firewall throughput (App-ID)	200 Gbps	120 Gbps	68 Gbps	68 Gbps	39 Gbps	18 Gbps
Threat Prevention throughput	100 Gbps	60 Gbps	30 Gbps	30 Gbps	20 Gbps	9 Gbps
IPsec VPN throughput	80 Gbps	48 Gbps	24 Gbps	24 Gbps	16 Gbps	8 Gbps
New sessions per second	1,200,000	720,000	462,000	462,000	348,000	171,000
Max sessions	40,000,000/80,000,000[3]	24,000,000/48,000,000[3]	64,000,000	32,000,000	8,000,000	4,000,000
Virtual systems (base/max[2])	25/225	25/225	25/225	25/225	25/125	10/20
Hardware Specifications	**PA-7080 System**	**PA-7050 System**	**PA-5280**	**PA-5260**	**PA-5250**	**PA-5220**
Interfaces supported NPC option 1[4]	Up to (20) QSFP+, (120) SFP+	Up to (12) QSFP+, (72) SFP+	(4) 100/1000/10G Cu, (16) 1G/10G SFP/SFP+, (4) 40G/100G QSFP28			(4) 100/1000/10G Cu, (16) 1G/10G SFP/SFP+, (4) 40G QSFP+
Management I/O	(2) 10/100/1000, (2) QSFP+ High availability, (1) 10/100/1000 Out-of-band management, (1) RJ45 console		(2) 10/100/1000 Cu, (1) 10/100/1000 Out-of-band management, (1) RJ45 console			
			(1) 40G/100G QSFP28 HA			(1) 40G QSFP+ HA
Rack mountable?	19U, 19" standard rack	9U, 19" standard rack or 14U, 19" standard rack with optional PAN-AIRDUCT kit	3U, 19" Standard rack			
Power supply	4x2500W AC (2400W/2700) expandable to 8	4x2500W AC (2400W/2700W)	2x1200W AC or DC (1:1 Fully redundant)			
Redundant power supply?	Yes		Yes			
Disk drives	2TB RAID1		System: 240GB SSD, RAID1. Log: 2TB HDD, RAID1			
Hot swap fans	Yes		Yes			

Performance and Capacities[1]	PA-5060	PA-5050	PA-5020	PA-3260	PA-3250	PA-3220	PA-3060	PA-3050	PA-3020
Firewall throughput (App-ID)	20 Gbps	10 Gbps	5 Gbps	8.8 Gbps	6.3 Gbps	5 Gbps	4 Gbps	4 Gbps	2 Gbps
Threat Prevention throughput	10 Gbps	5 Gbps	2 Gbps	4.7 Gbps	3 Gbps	2.2 Gbps	2 Gbps	2 Gbps	1 Gbps
IPsec VPN throughput	4 Gbps	4 Gbps	2 Gbps	4.8 Gbps	3.2 Gbps	2.5 Gbps	500 Mbps	500 Mbps	500 Mbps
New sessions per second	120,000	120,000	120,000	135,000	94,000	58,000	50,000	50,000	50,000
Max sessions	4,000,000	2,000,000	1,000,000	3,000,000	2,000,000	1,000,000	500,000	500,000	250,000
Virtual systems (base/max[2])	25/225	25/125	10/20	1/6	1/6	1/6	1/6	1/6	1/6
Hardware Specifications	**PA-5060**	**PA-5050**	**PA-5020**	**PA-3260**	**PA-3250**	**PA-3220**	**PA-3060**	**PA-3050**	**PA-3020**
Interfaces supported[4]	(12) 10/100/1000, (8) SFP, (4) 10 SFP+		12) 10/100/1000, (8) SFP	(12) 10/100/1000, (8) 1G/10G SFP/ SFP+, (4) 40G QSFP+	(12) 10/100/1000, (8) 1G/10G SFP/ SFP+	(12) 10/100/1000, (4) 1G SFP, (4) 1G/10G SFP/ SFP+	(8) 10/100/1000, (8) SFP, (2) 10 SFP+	(12) 10/100/1000, (8) SFP	
Management I/O	(2) 10/100/1000 High availability, (1) 10/100/1000 Out-of-band management, (1) RJ45 console			(1) 10/100/1000 out-of-band management port, (2) 10/100/1000 high availability, (1) 10G SFP+ high availability, (1) RJ-45 console port, (1) Micro USB			1) 10/100/1000 Out-of-band management, (2) 10/100/1000 High availability, (1) RJ-45 console		
Rack mountable?	2U, 19" Standard rack			2U, 19" standard rack (3.5" H x 20.53" D x 17.34" W)			1.5U, 19" Standard rack	1U, 19" Standard rack	
Power supply	Redundant 450W AC or DC			650-watt AC or DC (180/240)			Redundant 400W AC	250W AC	
Redundant power supply?	Yes			Yes			Yes	No	
Disk drives	120GB or 240GB SSD, RAID optional			240GB SSD			120GB SSD		
Hot swap fans	Yes			Yes			No		

Palo Alto Networks Platform Specifications and Features Summary

Performance and Capacities[1]	PA-850	PA-820	PA-500	PA-220	PA-200
Firewall throughput (App-ID)	1.9 Gbps	940 Mbps	250 Mbps	500 Mbps	100 Mbps
Threat Prevention throughput	780 Mbps	610 Mbps	100 Mbps	150 Mbps	50 Mbps
IPsec VPN throughput	500 Mbps	400 Mbps	50 Mbps	100 Mbps	50 Mbps
New sessions per second	9,500	8,300	7,500	4,200	1,000
Max sessions	192,000	128,000	64,000	64,000	64,000
Virtual systems (base)	1	1	N/A	1	N/A
Hardware Specifications	**PA-850**	**PA-820**	**PA-500**	**PA-220**	**PA-200**
Interfaces supported[4]	(4) 10/100/1000, (4/8) SFP, (0/4) 10 SFP+	(4) 10/100/1000, (8) SFP	(8) 10/100/1000	(8) 10/100/1000	(4) 10/100/1000
Management I/O	(1) 10/100/1000 Out-of-band management, (2) 10/100/1000 High availability, (1) RJ-45 console, (1) USB, (1) Micro USB console		(1) 10/100/1000 out-of-band management, (1) RJ-45 Console	(1) 10/100/1000 Out-of-band management, (1) RJ-45 Console, (1) USB, (1) Micro USB console	(1) 10/100/1000 Out-of-band management, (1) RJ-45 Console
Rack mountable?	1U, 19" Standard rack		1U, 19" standard rack	1.62"H X 6.29"D X 8.07"W	1.75" H x 7"D x 9.25"W
Power supply	Two 500W AC; One is redundant	200W	180W	Dual redundant 40W	40W
Redundant power supply?	Yes	No	No	Yes (optional)	No
Disk drives	240GB SSD		160GB	32GB EMMC	16GB SSD
Hot swap fans	No		No	No	No

Performance and Capacities[1]	VM-50/VM-50 Lite	VM-100/VM-200	VM-300/VM-1000-HV	VM-500	VM-700
Firewall throughput (App-ID)	200 Mbps	2 Gbps	4 Gbps	8 Gbps	16 Gbps
Threat Prevention throughput	100 Mbps	1 Gbps	2 Gbps	4 Gbps	8 Gbps
IPsec VPN throughput	100 Mbps	1 Gbps	1.8 Gbps	4 Gbps	6 Gbps
New sessions per second[1]	3,000	15,000	30,000	60,000	120,000
CPU configurations supported	2[6]	2	2,4	2,4,8	2,4,8,16
Dedicated memory (minimum)	4.0[8]/4.5GB	6.5GB	9GB	16GB	56GB
Dedicated disk drive capacity (minimum)	32GB[7]	60GB	60GB	60GB	60GB
VM-Series Supported Environments					
° VMware ESXi 5.1/5.5/6.0 (stand-alone) ° KVM on CentOS/RHEL and Ubuntu ° Microsoft Hyper-V (Windows 2012 R2 server)	Yes				
NSX Manager 6.0/6.1/6.2	No	Yes		Yes	No
Citrix Xen Server on SDX 10.1				No	
Amazon AWS		Yes (BYOL Only)	Yes (BYOL and Marketplace)	Yes (BYOL Only)	
Microsoft Azure					
Google Cloud					

(1) For VM-Series, performance and capacities may vary based on underlying virtualization infrastructure (hypervisor/cloud). Refer to the individual datasheets for detailed performance and testing information. (2) Adding virtual systems to the base quantity requires a separately purchased license. (3) Max session capacity for PA-7000 NPCs with standard memory/extended memory. (4) Optical/Copper transceivers are sold separately. (6) CPU oversubscription supported with up to 5 instances running on a 2 CPU configuration. (7) 60GB required at initial boot. VM-Series will use 32GB after license activation. (8) Supported with VM-50 Lite model only.

Palo Alto Networks Platform Specifications and Features Summary

Key Features	Supported Across All Platforms
Firewall	
Thousands of applications for visibility and control; ability to create custom applications; ability to manage unknown traffic based on policy	✓
User identification and control: VPNs, WLAN controllers, captive portal, proxies, Active Directory, eDirectory, Exchange, terminal services, syslog parsing, XML API	✓
Granular SSL decryption and inspection (inbound and outbound); per-policy SSH control (inbound and outbound)	✓
Networking: dynamic routing (RIP, OSPF, BGP, multiprotocol BGP), DHCP, DNS, NAT, route redistribution, ECMP, LLDP, BFD, tunnel content inspection	✓
QoS: policy-based traffic shaping (priority, guaranteed, maximum) per application, per user, per tunnel, based on DSCP classification	✓
Virtual systems: logical, separately managed firewall instances within a single physical firewall, with each virtual system's traffic kept separate	✓
Zone-based network segmentation and zone protection; DoS protection against flooding of new sessions	✓
Threat Prevention (subscription required)	
In-line malware prevention is automatically enforced through our payload-based signatures, which are updated daily	✓
Vulnerability-based protections block exploits and evasive techniques on network and application layers, including port scans, buffer overflows, packet fragmentation and obfuscation	✓
Command-and-control (C2) activity is stopped from exfiltrating data or delivering secondary malware payloads, while infected hosts are identified through DNS sinkholing	✓
URL Filtering (subscription required)	
Automatically prevent web-based attacks, including phishing links in emails, phishing sites, HTTP-based command-and-control, and pages that carry exploit kits	✓
Stops in-process credential phishing	✓
Custom URL categories, alerts and notification pages	✓
WildFire malware analysis (subscription required)	
Detects zero-day malware and exploits with layered, complimentary analysis techniquess	✓
Automates prevention within as few as five minutes across the network, endpoint and cloud	✓
Takes advantage of community-based data for protection, including over 22,000 subscribers	✓
AutoFocus threat intelligence (subscription required)	
Gain context and classification for attacks including malware family, adversary, and campaign to speed prioritization and response efforts	✓
Make security teams more effective with rich globally correlated threat analysis sourced from WildFire	✓
Use third party threat intelligence for automated prevention	✓
Magnifier behavioral analytics (subscription required)	
Automated profiling of user and device behavior through analysis of rich network, endpoint and cloud data stored in Palo Alto Networks Logging Service	✓
Accurate detection of stealthy network threats by identifying behavioral anomalies indicative of command and control, lateral movement, data exfiltration	✓
Integrated endpoint analysis to confirm attacks and accelerate investigations	✓
File and data filtering	
Bidirectional control over the unauthorized transfer of file types and Social Security numbers, credit card numbers and custom data patterns	✓
GlobalProtect mobile security (subscription required)	
Remote access VPN (SSL, IPSec, clientless) and mobile threat prevention and policy enforcement based on apps, users, content, device and device state	✓
BYOD: app-level VPN for user privacy	✓
Panorama management, logging and reporting (subscription required for managing multiple firewalls)	
Intuitive policy control with applications, users, threats, advanced malware protection, URL, file types, data patterns – all in the same policy	✓
Actionable insight into traffic and threats with Application Command Center (ACC), fully customizable reporting	✓
Aggregated logging and event correlation	✓
Consistent management of all hardware and all VM-Series, role-based access control, logical and hierarchical device groups, and templates	✓
GUI, CLI, XML-based REST API	✓

PAN-OS® 8.0 Administrator's Guide

paloaltonetworks.com/documentation

techDOCS paloalto NETWORKS

Contact Information

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

About the Documentation

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

Copyright

Palo Alto Networks, Inc.
www.paloaltonetworks.com

© 2017-2018 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

Last Revised

October 25, 2018

Table of Contents

Getting Started 17
Integrate the Firewall into Your Management Network 19
Determine Your Management Strategy 19
Perform Initial Configuration 20
Set Up Network Access for External Services 23
Register the Firewall 27
Create a New Support Account and Register a Firewall 27
Register a Firewall 29
Activate Licenses and Subscriptions 31
Install Content and Software Updates 33
Segment Your Network Using Interfaces and Zones 36
Network Segmentation for a Reduced Attack Surface 36
Configure Interfaces and Zones 36
Set Up a Basic Security Policy 39
Assess Network Traffic 42
Enable Basic WildFire Forwarding 44
Control Access to Web Content 46
Enable AutoFocus Threat Intelligence 48
Best Practices for Completing the Firewall Deployment 49
Best Practices for Securing Administrative Access 50

Firewall Administration 57
Management Interfaces 59
Use the Web Interface 60
Launch the Web Interface 60
Configure Banners, Message of the Day, and Logos 60
Use the Administrator Login Activity Indicators to Detect Account Misuse 62
Manage and Monitor Administrative Tasks 64
Commit, Validate, and Preview Firewall Configuration Changes 64
Use Global Find to Search the Firewall or Panorama Management Server 66
Manage Locks for Restricting Configuration Changes 67
Manage Configuration Backups 69
Save and Export Firewall Configurations 69
Revert Firewall Configuration Changes 70
Manage Firewall Administrators 73
Administrative Role Types 73
Configure an Admin Role Profile 74
Administrative Authentication 74
Configure Administrative Accounts and Authentication 75
Reference: Web Interface Administrator Access 81
Web Interface Access Privileges 81
Panorama Web Interface Access Privileges 135
Reference: Port Number Usage 139
Ports Used for Management Functions 139
Ports Used for HA 140
Ports Used for Panorama 141
Ports Used for GlobalProtect 142
Ports Used for User-ID 143
Reset the Firewall to Factory Default Settings 145

Bootstrap the Firewall....146
USB Flash Drive Support....146
Sample init-cfg.txt Files....147
Prepare a USB Flash Drive for Bootstrapping a Firewall....148
Bootstrap a Firewall Using a USB Flash Drive....150

Authentication....153
Authentication Types....155
External Authentication Services....155
Multi-Factor Authentication....155
SAML....156
Kerberos....157
TACACS+....158
RADIUS....158
LDAP....160
Local Authentication....160
Plan Your Authentication Deployment....161
Configure Multi-Factor Authentication....162
Configure MFA Between RSA SecurID and the Firewall....165
Configure MFA Between Okta and the Firewall....171
Configure MFA Between Duo and the Firewall....183
Configure SAML Authentication....191
Configure Kerberos Single Sign-On....195
Configure Kerberos Server Authentication....196
Configure TACACS+ Authentication....197
Configure RADIUS Authentication....199
Configure LDAP Authentication....202
Connection Timeouts for Authentication Servers....203
Guidelines for Setting Authentication Server Timeouts....203
Modify the PAN-OS Web Server Timeout....204
Modify the Captive Portal Session Timeout....204
Configure Local Database Authentication....205
Configure an Authentication Profile and Sequence....206
Test Authentication Server Connectivity....209
Authentication Policy....211
Authentication Timestamps....211
Configure Authentication Policy....212
Troubleshoot Authentication Issues....215

Certificate Management....217
Keys and Certificates....219
Certificate Revocation....222
Certificate Revocation List (CRL)....222
Online Certificate Status Protocol (OCSP)....223
Certificate Deployment....224
Set Up Verification for Certificate Revocation Status....225
Configure an OCSP Responder....225
Configure Revocation Status Verification of Certificates....226
Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption....226
Configure the Master Key....228
Obtain Certificates....229
Create a Self-Signed Root CA Certificate....229

Generate a Certificate 229
Import a Certificate and Private Key 231
Obtain a Certificate from an External CA 232
Export a Certificate and Private Key 234
Configure a Certificate Profile 235
Configure an SSL/TLS Service Profile 237
Replace the Certificate for Inbound Management Traffic 238
Configure the Key Size for SSL Forward Proxy Server Certificates 239
Revoke and Renew Certificates 240
Revoke a Certificate 240
Renew a Certificate 240
Secure Keys with a Hardware Security Module 241
Set up Connectivity with an HSM 241
Encrypt a Master Key Using an HSM 245
Store Private Keys on an HSM 246
Manage the HSM Deployment 247

High Availability 249
HA Overview 251
HA Concepts 252
HA Modes 252
HA Links and Backup Links 253
Device Priority and Preemption 256
Failover 256
LACP and LLDP Pre-Negotiation for Active/Passive HA 257
Floating IP Address and Virtual MAC Address 257
ARP Load-Sharing 259
Route-Based Redundancy 260
HA Timers 260
Session Owner 263
Session Setup 263
NAT in Active/Active HA Mode 265
ECMP in Active/Active HA Mode 266
Set Up Active/Passive HA 267
Prerequisites for Active/Passive HA 267
Configuration Guidelines for Active/Passive HA 267
Configure Active/Passive HA 269
Define HA Failover Conditions 274
Verify Failover 275
Set Up Active/Active HA 276
Prerequisites for Active/Active HA 276
Configure Active/Active HA 277
Determine Your Active/Active Use Case 281
HA Firewall States 295
Reference: HA Synchronization 297
What Settings Don't Sync in Active/Passive HA? 297
What Settings Don't Sync in Active/Active HA? 299
Synchronization of System Runtime Information 302

Monitoring 305
Use the Dashboard 307
Use the Application Command Center 309
ACC—First Look 309

ACC Tabs....311
ACC Widgets....312
Widget Descriptions....314
ACC Filters....318
Interact with the ACC....319
Use Case: ACC—Path of Information Discovery....322
Use the App Scope Reports....328
Summary Report....328
Change Monitor Report....329
Threat Monitor Report....330
Threat Map Report....331
Network Monitor Report....331
Traffic Map Report....332
Use the Automated Correlation Engine....334
Automated Correlation Engine Concepts....334
View the Correlated Objects....335
Interpret Correlated Events....335
Use the Compromised Hosts Widget in the ACC....337
Take Packet Captures....339
Types of Packet Captures....339
Disable Hardware Offload....339
Take a Custom Packet Capture....340
Take a Threat Packet Capture....344
Take an Application Packet Capture....345
Take a Packet Capture on the Management Interface....348
Monitor Applications and Threats....350
View and Manage Logs....351
Log Types and Severity Levels....351
View Logs....356
Filter Logs....357
Export Logs....357
Configure Log Storage Quotas and Expiration Periods....358
Schedule Log Exports to an SCP or FTP Server....358
Monitor Block List....360
View and Manage Reports....361
Report Types....361
View Reports....361
Configure the Expiration Period and Run Time for Reports....362
Disable Predefined Reports....362
Custom Reports....363
Generate Custom Reports....365
Generate Botnet Reports....367
Generate the SaaS Application Usage Report....368
Manage PDF Summary Reports....371
Generate User/Group Activity Reports....373
Manage Report Groups....374
Schedule Reports for Email Delivery....375
Use External Services for Monitoring....377
Configure Log Forwarding....378
Configure Email Alerts....381
Use Syslog for Monitoring....382
Configure Syslog Monitoring....382
Syslog Field Descriptions....384
SNMP Monitoring and Traps....418
SNMP Support....418

Use an SNMP Manager to Explore MIBs and Objects....419
Enable SNMP Services for Firewall-Secured Network Elements....421
Monitor Statistics Using SNMP....422
Forward Traps to an SNMP Manager....423
Supported MIBs....425
Forward Logs to an HTTP(S) Destination....433
NetFlow Monitoring....436
Configure NetFlow Exports....436
NetFlow Templates....437
Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors....443

User-ID....445
User-ID Overview....447
User-ID Concepts....449
Group Mapping....449
User Mapping....449
Enable User-ID....453
Map Users to Groups....456
Map IP Addresses to Users....459
Create a Dedicated Service Account for the User-ID Agent....459
Configure User Mapping Using the Windows User-ID Agent....461
Configure User Mapping Using the PAN-OS Integrated User-ID Agent....472
Configure User-ID to Monitor Syslog Senders for User Mapping....474
Map IP Addresses to Usernames Using Captive Portal....482
Configure User Mapping for Terminal Server Users....487
Send User Mappings to User-ID Using the XML API....495
Enable User- and Group-Based Policy....496
Enable Policy for Users with Multiple Accounts....497
Verify the User-ID Configuration....499
Deploy User-ID in a Large-Scale Network....501
Deploy User-ID for Numerous Mapping Information Sources....501
Redistribute User Mappings and Authentication Timestamps....505

App-ID....509
App-ID Overview....511
Manage Custom or Unknown Applications....512
Manage New App-IDs Introduced in Content Releases....513
Review New App-IDs....513
Review New App-IDs Since Last Content Version....513
Review New App-ID Impact on Existing Policy Rules....514
Disable or Enable App-IDs....515
Prepare Policy Updates for Pending App-IDs....516
Use Application Objects in Policy....518
Create an Application Group....518
Create an Application Filter....518
Create a Custom Application....519
Applications with Implicit Support....523
Application Level Gateways....526
Disable the SIP Application-level Gateway (ALG)....528

Threat Prevention....531
Set Up Antivirus, Anti-Spyware, and Vulnerability Protection....533

Create Threat Exceptions....536
Set Up Data Filtering....538
Set Up File Blocking....540
Prevent Brute Force Attacks....543
Customize the Action and Trigger Conditions for a Brute Force Signature....544
Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions....547
Enable Evasion Signatures....553
Best Practices for Applications and Threats Content Updates....554
Prevent Credential Phishing....557
Methods to Check for Corporate Credential Submissions....557
Configure Credential Detection with the Windows-based User-ID Agent....559
Set Up Credential Phishing Prevention....561
Share Threat Intelligence with Palo Alto Networks....564
What Telemetry Data Does the Firewall Collect?....564
Passive DNS Monitoring....565
Enable Telemetry....566
Use DNS Queries to Identify Infected Hosts on the Network....568
DNS Sinkholing....568
Configure DNS Sinkholing for a List of Custom Domains....569
Configure the Sinkhole IP Address to a Local Server on Your Network....571
Identify Infected Hosts....573
Monitor Blocked IP Addresses....576
Learn More About and Assess Threats....578
Assess Firewall Artifacts with AutoFocus....578
Learn More About Threat Signatures....582
Monitor Activity and Create Custom Reports Based on Threat Categories....584
Content Delivery Network Infrastructure for Dynamic Updates....586
Threat Prevention Resources....587

Decryption....589
Decryption Overview....591
Decryption Concepts....592
Keys and Certificates for Decryption Policies....592
SSL Forward Proxy....594
SSL Forward Proxy Decryption Profile....595
SSL Inbound Inspection....597
SSL Inbound Inspection Decryption Profile....598
SSL Protocol Settings Decryption Profile....599
SSH Proxy....601
SSH Proxy Decryption Profile....602
Decryption Profile for No Decryption....603
Decryption Mirroring....604
SSL Decryption for Elliptical Curve Cryptography (ECC) Certificates....604
Perfect Forward Secrecy (PFS) Support for SSL Decryption....605
High Availability (HA) Support for Decrypted Sessions....605
Prepare to Deploy Decryption....607
Work with Stakeholders to Develop a Decryption Deployment Strategy....607
Develop a PKI Rollout Plan....609
Size the Firewall Decryption Deployment....610
Plan a Staged, Prioritized Deployment....611
Define Traffic to Decrypt....613
Create a Decryption Profile....613
Create a Decryption Policy Rule....615
Configure SSL Forward Proxy....618

Configure SSL Inbound Inspection....622
Configure SSH Proxy....624
Configure Server Certificate Verification for Undecrypted Traffic....625
Decryption Exclusions....626
Palo Alto Networks Predefined Decryption Exclusions....626
Exclude a Server from Decryption for Technical Reasons....627
Create a Policy-Based Decryption Exclusion....628
Enable Users to Opt Out of SSL Decryption....631
Configure Decryption Port Mirroring....633
Verify Decryption....636
Temporarily Disable SSL Decryption....639

URL Filtering....641
URL Filtering Overview....643
URL Filtering Vendors....643
Interaction Between App-ID and URL Categories....643
PAN-DB Private Cloud....644
URL Filtering Concepts....646
URL Categories....646
URL Filtering Profile....647
URL Filtering Profile Actions....647
URL Category Exception Lists....649
External Dynamic List for URLs....651
Container Pages....651
HTTP Header Logging....652
URL Filtering Response Pages....652
URL Category as Policy Match Criteria....655
PAN-DB Categorization....657
Enable a URL Filtering Vendor....659
Enable PAN-DB URL Filtering....659
Enable BrightCloud URL Filtering....660
Determine URL Filtering Policy Requirements....662
Configure URL Filtering....664
Use an External Dynamic List in a URL Filtering Profile....667
Customize the URL Filtering Response Pages....669
Allow Password Access to Certain Sites....670
Safe Search Enforcement....672
Safe Search Settings for Search Providers....672
Block Search Results when Strict Safe Search is not Enabled....673
Transparently Enable Safe Search for Users....676
Monitor Web Activity....680
Monitor Web Activity of Network Users....680
View the User Activity Report....682
Configure Custom URL Filtering Reports....683
Set Up the PAN-DB Private Cloud....686
Configure the PAN-DB Private Cloud....686
Configure the Firewalls to Access the PAN-DB Private Cloud....689
URL Filtering Use Cases....691
Use Case: Control Web Access....691
Use Case: Use URL Categories for Policy Matching....694
Troubleshoot URL Filtering....696
Problems Activating PAN-DB....696
PAN-DB Cloud Connectivity Issues....696
URLs Classified as Not-Resolved....698

Incorrect Categorization....698
URL Database Out of Date....699

Quality of Service....701
QoS Overview....703
QoS Concepts....705
QoS for Applications and Users....705
QoS Policy....705
QoS Profile....705
QoS Classes....706
QoS Priority Queuing....706
QoS Bandwidth Management....706
QoS Egress Interface....707
QoS for Clear Text and Tunneled Traffic....708
Configure QoS....709
Configure QoS for a Virtual System....714
Enforce QoS Based on DSCP Classification....719
QoS Use Cases....721
Use Case: QoS for a Single User....721
Use Case: QoS for Voice and Video Applications....723

VPNs....727
VPN Deployments....729
Site-to-Site VPN Overview....730
Site-to-Site VPN Concepts....731
IKE Gateway....731
Tunnel Interface....731
Tunnel Monitoring....732
Internet Key Exchange (IKE) for VPN....732
IKEv2....734
Set Up Site-to-Site VPN....738
Set Up an IKE Gateway....738
Define Cryptographic Profiles....743
Set Up an IPSec Tunnel....745
Set Up Tunnel Monitoring....747
Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel....748
Test VPN Connectivity....750
Interpret VPN Error Messages....751
Site-to-Site VPN Quick Configs....752
Site-to-Site VPN with Static Routing....752
Site-to-Site VPN with OSPF....755
Site-to-Site VPN with Static and Dynamic Routing....758

Large Scale VPN (LSVPN)....765
LSVPN Overview....767
Create Interfaces and Zones for the LSVPN....768
Enable SSL Between GlobalProtect LSVPN Components....770
About Certificate Deployment....770
Deploy Server Certificates to the GlobalProtect LSVPN Components....770
Deploy Client Certificates to the GlobalProtect Satellites Using SCEP....772
Configure the Portal to Authenticate Satellites....775
Configure GlobalProtect Gateways for LSVPN....777

Configure the GlobalProtect Portal for LSVPN 780
GlobalProtect Portal for LSVPN Prerequisite Tasks 780
Configure the Portal 780
Define the Satellite Configurations 781
Prepare the Satellite to Join the LSVPN 785
Verify the LSVPN Configuration 787
LSVPN Quick Configs 788
Basic LSVPN Configuration with Static Routing 788
Advanced LSVPN Configuration with Dynamic Routing 790
Advanced LSVPN Configuration with iBGP 792

Networking 797
Configure Interfaces 799
Tap Interfaces 799
Virtual Wire Interfaces 799
Layer 2 Interfaces 806
Layer 3 Interfaces 809
Configure an Aggregate Interface Group 818
Use Interface Management Profiles to Restrict Access 820
Virtual Routers 822
Service Routes 824
Static Routes 826
Static Route Overview 826
Static Route Removal Based on Path Monitoring 826
Configure a Static Route 828
Configure Path Monitoring for a Static Route 829
RIP 832
OSPF 834
OSPF Concepts 834
Configure OSPF 836
Configure OSPFv3 838
Configure OSPF Graceful Restart 840
Confirm OSPF Operation 841
BGP 843
BGP Overview 843
MP-BGP 843
Configure BGP 845
Configure a BGP Peer with MP-BGP for IPv4 or IPv6 Unicast 849
Configure a BGP Peer with MP-BGP for IPv4 Multicast 851
BGP Confederations 852
IP Multicast 857
IGMP 857
PIM 858
Configure IP Multicast 863
View IP Multicast Information 868
Route Redistribution 870
DHCP 873
DHCP Overview 873
Firewall as a DHCP Server and Client 874
DHCP Messages 874
DHCP Addressing 875
DHCP Options 876
Configure an Interface as a DHCP Server 878
Configure an Interface as a DHCP Client 881

Configure the Management Interface as a DHCP Client....882
Configure an Interface as a DHCP Relay Agent....884
Monitor and Troubleshoot DHCP....884
DNS....887
DNS Overview....887
DNS Proxy Object....888
DNS Server Profile....888
Multi-Tenant DNS Deployments....889
Configure a DNS Proxy Object....890
Configure a DNS Server Profile....891
Use Case 1: Firewall Requires DNS Resolution for Management Purposes....892
Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System....893
Use Case 3: Firewall Acts as DNS Proxy Between Client and Server....895
DNS Proxy Rule and FQDN Matching....896
NAT....900
NAT Policy Rules....900
Source NAT and Destination NAT....902
NAT Rule Capacities....904
Dynamic IP and Port NAT Oversubscription....904
Dataplane NAT Memory Statistics....906
Configure NAT....907
NAT Configuration Examples....912
NPTv6....919
NPTv6 Overview....919
How NPTv6 Works....920
NDP Proxy....922
NPTv6 and NDP Proxy Example....923
Create an NPTv6 Policy....924
NAT64....927
NAT64 Overview....927
IPv4-Embedded IPv6 Address....928
DNS64 Server....928
Path MTU Discovery....928
IPv6-Initiated Communication....929
Configure NAT64 for IPv6-Initiated Communication....930
Configure NAT64 for IPv4-Initiated Communication....932
Configure NAT64 for IPv4-Initiated Communication with Port Translation....934
ECMP....937
ECMP Load-Balancing Algorithms....937
ECMP Model, Interface, and IP Routing Support....938
Configure ECMP on a Virtual Router....938
Enable ECMP for Multiple BGP Autonomous Systems....940
Verify ECMP....941
LLDP....942
LLDP Overview....942
Supported TLVs in LLDP....943
LLDP Syslog Messages and SNMP Traps....944
Configure LLDP....944
View LLDP Settings and Status....946
Clear LLDP Statistics....947
BFD....948
BFD Overview....948
Configure BFD....950
Reference: BFD Details....955

Session Settings and Timeouts ... 959
Transport Layer Sessions ... 959
TCP ... 959
UDP ... 963
ICMP ... 964
Control Specific ICMP or ICMPv6 Types and Codes ... 965
Configure Session Timeouts ... 966
Configure Session Settings ... 967
Session Distribution Policies ... 970
Prevent TCP Split Handshake Session Establishment ... 973
Tunnel Content Inspection ... 975
Tunnel Content Inspection Overview ... 975
Configure Tunnel Content Inspection ... 978
View Inspected Tunnel Activity ... 982
View Tunnel Information in Logs ... 983
Create a Custom Report Based on Tagged Tunnel Traffic ... 983

Policy ... 985
Policy Types ... 987
Security Policy ... 988
Components of a Security Policy Rule ... 988
Security Policy Actions ... 990
Create a Security Policy Rule ... 991
Policy Objects ... 994
Security Profiles ... 995
Create a Security Profile Group ... 1000
Set Up or Override a Default Security Profile Group ... 1001
Enumeration of Rules Within a Rulebase ... 1003
Move or Clone a Policy Rule or Object to a Different Virtual System ... 1005
Use Tags to Group and Visually Distinguish Objects ... 1006
Create and Apply Tags ... 1006
Modify Tags ... 1007
Use the Tag Browser ... 1007
Use an External Dynamic List in Policy ... 1011
External Dynamic List ... 1011
Formatting Guidelines for an External Dynamic List ... 1012
Palo Alto Networks Malicious IP Address Feeds ... 1013
Configure the Firewall to Access an External Dynamic List ... 1014
Retrieve an External Dynamic List from the Web Server ... 1016
View External Dynamic List Entries ... 1016
Exclude Entries from an External Dynamic List ... 1017
Enforce Policy on an External Dynamic List ... 1017
Find External Dynamic Lists That Failed Authentication ... 1019
Disable Authentication for an External Dynamic List ... 1020
Register IP Addresses and Tags Dynamically ... 1021
Monitor Changes in the Virtual Environment ... 1022
Enable VM Monitoring to Track Changes on the Virtual Network ... 1022
Attributes Monitored on Virtual Machines in Cloud Platforms ... 1024
Use Dynamic Address Groups in Policy ... 1024
CLI Commands for Dynamic IP Addresses and Tags ... 1028
Identify Users Connected through a Proxy Server ... 1030
Use XFF Values for Policies and Logging Source Users ... 1030
Use the IP Address in the XFF Header to Troubleshoot Events ... 1031
Policy-Based Forwarding ... 1033

PBF....1033
Create a Policy-Based Forwarding Rule....1035
Use Case: PBF for Outbound Access with Dual ISPs....1036

Virtual Systems....1043
Virtual Systems Overview....1045
Virtual System Components and Segmentation....1045
Benefits of Virtual Systems....1046
Use Cases for Virtual Systems....1046
Platform Support and Licensing for Virtual Systems....1046
Administrative Roles for Virtual Systems....1047
Shared Objects for Virtual Systems....1047
Communication Between Virtual Systems....1048
Inter-VSYS Traffic That Must Leave the Firewall....1048
Inter-VSYS Traffic That Remains Within the Firewall....1048
Inter-VSYS Communication Uses Two Sessions....1050
Shared Gateway....1051
External Zones and Shared Gateway....1051
Networking Considerations for a Shared Gateway....1052
Configure Virtual Systems....1053
Configure Inter-Virtual System Communication within the Firewall....1056
Configure a Shared Gateway....1057
Customize Service Routes for a Virtual System....1058
Customize Service Routes to Services for Virtual Systems....1058
Configure a PA-7000 Series Firewall for Logging Per Virtual System....1059
Configure Administrative Access Per Virtual System or Firewall....1060
Virtual System Functionality with Other Features....1062

Zone Protection and DoS Protection....1063
Network Segmentation Using Zones....1065
How Do Zones Protect the Network?....1066
Zone Defense....1067
Zone Defense Tools....1067
How Do the Zone Defense Tools Work?....1068
Zone Protection Profiles....1069
Packet Buffer Protection....1071
DoS Protection Profiles and Policy Rules....1071
Configure Zone Protection to Increase Network Security....1074
Configure Reconnaissance Protection....1074
Configure Packet Based Attack Protection....1074
Configure Protocol Protection....1075
Configure Packet Buffer Protection....1079
DoS Protection Against Flooding of New Sessions....1081
Multiple-Session DoS Attack....1081
Single-Session DoS Attack....1084
Configure DoS Protection Against Flooding of New Sessions....1084
End a Single Session DoS Attack....1086
Identify Sessions That Use an Excessive Percentage of the Packet Buffer....1087
Discard a Session Without a Commit....1089

Certifications....1091
Enable FIPS and Common Criteria Support....1093

Access the Maintenance Recovery Tool (MRT)..... 1093
Change the Operational Mode to FIPS-CC Mode..... 1094
FIPS-CC Security Functions..... 1096

Getting Started

The following topics provide detailed steps to help you deploy a new Palo Alto Networks next-generation firewall. They provide details for integrating a new firewall into your network, registering the firewall, activating licenses and subscriptions, and configuring basic security policies and threat prevention features.

After you perform the basic configuration steps required to integrate the firewall into your network, you can use the rest of the topics in this guide to help you deploy the comprehensive security platform features as necessary to address your network security needs.

> Integrate the Firewall into Your Management Network
> Register the Firewall
> Activate Licenses and Subscriptions
> Install Content and Software Updates
> Segment Your Network Using Interfaces and Zones
> Set Up a Basic Security Policy
> Assess Network Traffic
> Enable Basic WildFire Forwarding
> Control Access to Web Content
> Enable AutoFocus Threat Intelligence
> Best Practices for Completing the Firewall Deployment
> Best Practices for Securing Administrative Access

© 2018 Palo Alto Networks, Inc.

Integrate the Firewall into Your Management Network

All Palo Alto Networks firewalls provide an out-of-band management port (MGT) that you can use to perform the firewall administration functions. By using the MGT port, you separate the management functions of the firewall from the data processing functions, safeguarding access to the firewall and enhancing performance. When using the web interface, you must perform all initial configuration tasks from the MGT port even if you plan to use an in-band data port for managing your firewall going forward.

Some management tasks, such as retrieving licenses and updating the threat and application signatures on the firewall require access to the Internet. If you do not want to enable external access to your MGT port, you will need to either set up an in-band data port to provide access to required external services (using service routes) or plan to manually upload updates regularly.



Do not enable access to your management interface from the internet or from other untrusted zones inside your enterprise security boundary. This applies whether you use the dedicated management port (MGT) or you configured a data port as your management interface. When integrating the firewall into the management interface, follow the Best Practices for Securing Administrative Access *to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.*

The following topics describe how to perform the initial configuration steps that are necessary to integrate a new firewall into the management network and deploy it in a basic security configuration.

- Determine Your Management Strategy
- Perform Initial Configuration
- Set Up Network Access for External Services



The following topics describe how to integrate a single Palo Alto Networks next-generation firewall into your network. However, for redundancy, consider deploying a pair of firewalls in a High Availability *configuration.*

Determine Your Management Strategy

The Palo Alto Networks firewall can be configured and managed locally or it can be managed centrally using Panorama, the Palo Alto Networks centralized security management system. If you have six or more firewalls deployed in your network, use Panorama to achieve the following benefits:

- Reduce the complexity and administrative overhead in managing configuration, policies, software and dynamic content updates. Using device groups and templates on Panorama, you can effectively manage firewall-specific configuration locally on a firewall and enforce shared policies across all firewalls or device groups.
- Aggregate data from all managed firewalls and gain visibility across all the traffic on your network. The Application Command Center (ACC) on Panorama provides a single glass pane for unified reporting across all the firewalls, allowing you to centrally analyze, investigate and report on network traffic, security incidents and administrative modifications.

The procedures that follow describe how to manage the firewall using the local web interface. If you want to use Panorama for centralized management, first Perform Initial Configuration and verify that the firewall can establish a connection to Panorama. From that point on you can use Panorama to configure your firewall centrally.

© 2018 Palo Alto Networks, Inc.

Perform Initial Configuration

By default, the firewall has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other firewall configuration tasks. You must perform these initial configuration tasks either from the MGT port, even if you do not plan to use this port for your firewall management, or using a direct serial connection to the console port on the firewall.

Before performing initial configuration, review the Best Practices for Securing Administrative Access *to ensure that you securing firewall management access in a way that will allow you to detect and defend against vulnerability exploits that could lead to the infiltration and/or the loss of sensitive data.*

STEP 1 | Install your firewall and connect power to it.

If your firewall model has dual power supplies, connect the second power supply for redundancy. Refer to the hardware reference guide *for your model for details.*

STEP 2 | Gather the required information from your network administrator.

- IP address for MGT port
- Netmask
- Default gateway
- DNS server address

STEP 3 | Connect your computer to the firewall.

You can connect to the firewall in one of the following ways:

- Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1). Wait a few minutes for the boot-up sequence to complete; when the firewall is ready, the prompt changes to the name of the firewall, for example `PA-500 login`.
- Connect an RJ-45 Ethernet cable from your computer to the MGT port on the firewall. From a browser, go to **`https://192.168.1.1`**.

You may need to change the IP address on your computer to an address in the 192.168.1.0/24 network, such as 192.168.1.2, in order to access this URL.

STEP 4 | When prompted, log in to the firewall.

You must log in using the default username and password (admin/admin). The firewall will begin to initialize.

STEP 5 | Configure the MGT port.

1. Select **Device** > **Setup** > **Interfaces** and edit the **Management** interface.
2. Configure the address settings for the MGT interface using one of the following methods:
   - To configure static IP address settings for the MGT interface, set the **IP Type** to **Static** and enter the **IP Address**, **Netmask**, and **Default Gateway**.
   - To dynamically configure the MGT interface address settings, set the **IP Type** to **DHCP Client**. To use this method, you must Configure the Management Interface as a DHCP Client.

© 2018 Palo Alto Networks, Inc.

To prevent unauthorized access to the management interface, it is a best practice to Add the Permitted IP Addresses from which an administrator can access the MGT interface.

3. Set the **Speed** to **auto-negotiate**.
4. Select which management services to allow on the interface.

Make sure Telnet and HTTP are not selected because these services use plaintext and are not as secure as the other services and could compromise administrator credentials.

5. Click **OK**.

STEP 6 | Configure DNS, update server, and proxy server settings.

You must manually configure at least one DNS server on the firewall or it will not be able to resolve hostnames; it will not use DNS server settings from another source, such as an ISP.

1. Select **Device** > **Setup** > **Services**.
   - For multi-virtual system platforms, select **Global** and edit the Services section.
   - For single virtual system platforms, edit the Services section.
2. On the **Services** tab, for **DNS**, click one of the following:
   - **Servers**—Enter the **Primary DNS Server** address and **Secondary DNS Server** address.
   - **DNS Proxy Object**—From the drop-down, select the **DNS Proxy** that you want to use to configure global DNS services, or click **DNS Proxy** to configure a new DNS proxy object.
3. Click **OK**.

STEP 7 | Configure date and time (NTP) settings.

1. Select **Device** > **Setup** > **Services**.
   - For multi-virtual system platforms, select **Global** and edit the Services section.
   - For single virtual system platforms, edit the Services section.
2. On the **NTP** tab, to use the virtual cluster of time servers on the Internet, enter the hostname `pool.ntp.org` as the **Primary NTP Server** or enter the IP address of your primary NTP server.
3. (Optional) Enter a **Secondary NTP Server** address.
4. (Optional) To authenticate time updates from the NTP server(s), for **Authentication Type**, select one of the following for each server:
   - **None**—(Default) Disables NTP authentication.
   - **Symmetric Key**—Firewall uses symmetric key exchange (shared secrets) to authenticate time updates.
     - **Key ID**—Enter the Key ID (1-65534).
     - **Algorithm**—Select the algorithm to use in NTP authentication (**MD5** or **SHA1**).
   - **Autokey**—Firewall uses autokey (public key cryptography) to authenticate time updates.
5. Click **OK**.

STEP 8 | (Optional) Configure general firewall settings as needed.

1. Select **Device** > **Setup** > **Management** and edit the General Settings.
2. Enter a **Hostname** for the firewall and enter your network **Domain** name. The domain name is just a label; it will not be used to join the domain.
3. Enter **Login Banner** text that informs users who are about to log in that they require authorization to access the firewall management functions.

© 2018 Palo Alto Networks, Inc.

As a best practice, avoid using welcoming verbiage. Additionally, you should ask your legal department to review the banner message to ensure it adequately warns that unauthorized access is prohibited.

4. Enter the **Latitude** and **Longitude** to enable accurate placement of the firewall on the world map.
5. Click **OK**.

STEP 9 | Set a secure password for the admin account.

Be sure to configure the password complexity settings to ensure a strict password policy that requires frequent password changes.

1. Select **Device** > **Administrators**.
2. Select the **admin** role.
3. Enter the current default password and the new password.
4. Click **OK** to save your settings.

STEP 10 | Commit your changes.

When the configuration changes are saved, you lose connectivity to the web interface because the IP address has changed.

Click **Commit** at the top right of the web interface. The firewall can take up to 90 seconds to save your changes.

STEP 11 | Connect the firewall to your network.

1. Disconnect the firewall from your computer.
2. Connect the MGT port to a switch port on your management network using an RJ-45 Ethernet cable. Make sure that the switch port you cable the firewall to is configured for auto-negotiation.

STEP 12 | Open an SSH management session to the firewall.

Using a terminal emulation software, such as PuTTY, launch an SSH session to the firewall using the new IP address you assigned to it.

STEP 13 | Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server.

You can do this in one of the following ways:

- (Recommended) If you do not want to allow external network access to the MGT interface, you will need to set up a data port to retrieve required service updates. Continue to Set Up Network Access for External Services.
- If you do plan to allow external network access from the MGT interface, verify that you have connectivity and then proceed to Register the Firewall and Activate Licenses and Subscriptions.

1. Use the ping utility to verify network connectivity to the Palo Alto Networks Update server as shown in the following example. Verify that DNS resolution occurs and the response includes the IP address for the Update server; the update server does not respond to a ping request.

```
admin@PA-220 > ping host updates.paloaltonetworks.com

PING updates.paloaltonetworks.com (10.101.16.13) 56(84) bytes of data.

From 192.168.1.1 icmp_seq=1 Destination Host Unreachable
```


© 2018 Palo Alto Networks, Inc.

```
From 192.168.1.1 icmp_seq=2 Destination Host Unreachable

From 192.168.1.1 icmp_seq=3 Destination Host Unreachable

From 192.168.1.1 icmp_seq=4 Destination Host Unreachable
```

After verifying DNS resolution, press Ctrl+C to stop the ping request.

2. Use the following CLI command to retrieve information on the support entitlement for the firewall from the Palo Alto Networks update server:

```
request support check
```

If you have connectivity, the update server will respond with the support status for your firewall. Because your firewall is not registered, the update server will return the following message:

```
Contact Us

https://www.paloaltonetworks.com/company/contact-us.html

Support Home

https://www.paloaltonetworks.com/support/tabs/overview.html

Device not found on this update server
```

Set Up Network Access for External Services

By default, the firewall uses the MGT port to access remote services, such as DNS servers, content updates, and license retrieval. If you do not want to enable external network access to your management network, you must set up an in-band data port to provide access to required external services and set up service routes to instruct the firewall what port to use to access the external services.

Do not enable management access from the internet or from other untrusted zones inside your enterprise security boundary. Follow the Best Practices for Securing Administrative Access to ensure that you are properly securing your firewall.

This task requires familiarity with firewall interfaces, zones, and policies. For more information on these topics, see Configure Interfaces and Zones and Set Up a Basic Security Policy.

STEP 1 | Decide which port you want to use for access to external services and connect it to your switch or router port.

The interface you use must have a static IP address.

STEP 2 | Log in to the web interface.

Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page.

STEP 3 | (Optional) The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and zones). If you

© 2018 Palo Alto Networks, Inc.

do not plan to use this virtual wire configuration, you must manually delete the configuration to prevent it from interfering with other interface settings you define.

You must delete the configuration in the following order:

1. To delete the default security policy, select **Policies** > **Security**, select the rule, and click **Delete**.
2. To delete the default virtual wire, select **Network** > **Virtual Wires**, select the virtual wire and click **Delete**.
3. To delete the default trust and untrust zones, select **Network** > **Zones**, select each zone and click **Delete**.
4. To delete the interface configurations, select **Network** > **Interfaces** and then select each interface (ethernet1/1 and ethernet1/2) and click **Delete**.
5. **Commit** the changes.

STEP 4 | Configure the interface you want the firewall to use for access to external services, such as DNS, NTP, Palo Alto Networks Services, or AutoFocus.

1. Select **Network** > **Interfaces** and select the interface that corresponds to the port you cabled in Step 1.
2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**.
3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**.
4. In the Zone dialog, enter a **Name** for new zone, for example Management, and then click **OK**.
5. Select the **IPv4** tab, select the **Static** radio button, and click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.254/24. You must use a static IP address on this interface.
6. Select **Advanced** > **Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**.
7. Enter a **Name** for the profile, such as allow_ping, and then select the services you want to allow on the interface. For the purposes of allowing access to the external services, you probably only need to enable **Ping** and then click **OK**.

   These services provide management access to the firewall, so only select the services that correspond to the management activities you want to allow on this interface. For example, if you plan to use the MGT interface for firewall configuration tasks through the web interface or CLI, you would not want to enable HTTP, HTTPS, SSH, or Telnet so that you could prevent unauthorized access through this interface (and if you did allow those services, you should limit access to a specific set of Permitted IP Addresses). For details, see Use Interface Management Profiles to Restrict Access. *Make sure you are following the* Best Practices for Securing Administrative Access.
8. To save the interface configuration, click **OK**.

STEP 5 | Configure the Service Routes.

By default, the firewall uses the MGT port to access the external services it requires. To change the interface the firewall uses to send requests to external services, you must edit the service routes.

This example shows how to set up global service routes. For information on setting up network access to external services on a virtual system basis rather than a global basis, see Customize Service Routes to Services for Virtual Systems.

1. Select **Device** > **Setup** > **Services** > **Global** and click **Service Route Configuration**.



© 2018 Palo Alto Networks, Inc.

For the purposes of activating your licenses and getting the most recent content and software updates, you will want to change the service route for DNS, Palo Alto Networks Services, URL Updates, and AutoFocus.

2. Click the **Customize** radio button, and select one of the following:
   - For a predefined service, select **IPv4** or **IPv6** and click the link for the service. To limit the drop-down list for Source Address, select **Source Interface** and select the interface you just configured. Then select a Source Address (from that interface) as the service route.

     If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you to select an IP address.
   - To create a service route for a custom destination, select **Destination**, and click **Add**. Enter a **Destination** IP address. An incoming packet with a destination address that matches this address will use as its source the Source Address you specify for this service route. To limit the drop-down for Source Address, select a **Source Interface**. If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you to select an IP address.
3. Click **OK** to save the settings.
4. Repeat steps b-c above for each service route you want to modify.
5. **Commit** your changes.

STEP 6 | Configure an external-facing interface and an associated zone and then create a security policy rule to allow the firewall to send service requests from the internal zone to the external zone.

1. Select **Network** > **Interfaces** and then select the external-facing interface. Select **Layer3** as the **Interface Type**, **Add** the **IP** address (on the **IPv4** or **IPv6** tab), and create the associated **Security Zone** (on the **Config** tab), such as Internet. This interface must have a static IP address; you do not need to set up management services on this interface.
2. To set up a security rule that allows traffic from your internal network to the Palo Alto Networks update server, select **Policies** > **Security** and click **Add**.

As a best practice when creating Security policy rules, use application-based rules instead of port-based rules to ensure that you are accurately identifying the underlying application regardless of the port, protocol, evasive tactics, or encryption in use. Always leave the Service set to application-default. In this case, create a security policy rule that allows access to the update server (and other Palo Alto Networks services).



STEP 7 | Create a NAT policy rule.

1. If you are using a private IP address on the internal-facing interface, you will need to create a source NAT rule to translate the address to a publicly routable address. Select **Policies** > **NAT** and then click **Add**. At a minimum you must define a name for the rule (**General** tab), specify a source and destination zone, Management to Internet in this case (**Original Packet** tab), and define the source address translation settings (**Translated Packet** tab) and then click **OK**.
2. **Commit** your changes.



STEP 8 | Verify that you have connectivity from the data interface to the external services, including the default gateway, and the Palo Alto Networks Update Server.

© 2018 Palo Alto Networks, Inc.

After you verify you have the required network connectivity, continue to Register the Firewall and Activate Licenses and Subscriptions.

1. Use the ping utility to verify network connectivity to the Palo Alto Networks Update server as shown in the following example. Verify that DNS resolution occurs and the response includes the IP address for the Update server; the update server does not respond to a ping request.

```
admin@PA-220 > ping host updates.paloaltonetworks.com
PING updates.paloaltonetworks.com (10.101.16.13) 56(84) bytes of data.
From 192.168.1.1 icmp_seq=1 Destination Host Unreachable
From 192.168.1.1 icmp_seq=2 Destination Host Unreachable
From 192.168.1.1 icmp_seq=3 Destination Host Unreachable
From 192.168.1.1 icmp_seq=4 Destination Host Unreachable
```



After verifying DNS resolution, press Ctrl+C to stop the ping request.

2. Use the following CLI command to retrieve information on the support entitlement for the firewall from the Palo Alto Networks update server:

```
request support check
```

If you have connectivity, the update server will respond with the support status for your firewall. Because your firewall is not registered, the update server will return the following message:

```
Contact Us
https://www.paloaltonetworks.com/company/contact-us.html
Support Home
https://www.paloaltonetworks.com/support/tabs/overview.html
Device not found on this update server
```


© 2018 Palo Alto Networks, Inc.

Register the Firewall

Before you can activate support and other licenses and subscriptions, you must first register the firewall. Before you can register a firewall, though, you must first have an active support account. Perform one of the following tasks depending on whether you have an active support account:

- If you don't have an active support account, then Create a New Support Account and Register a Firewall on page 27.
- If you already have an active support account, then you are ready to Register a Firewall on page 29.

If you are registering a VM-Series firewall, refer to the VM-Series Deployment Guide for instructions.

Create a New Support Account and Register a Firewall

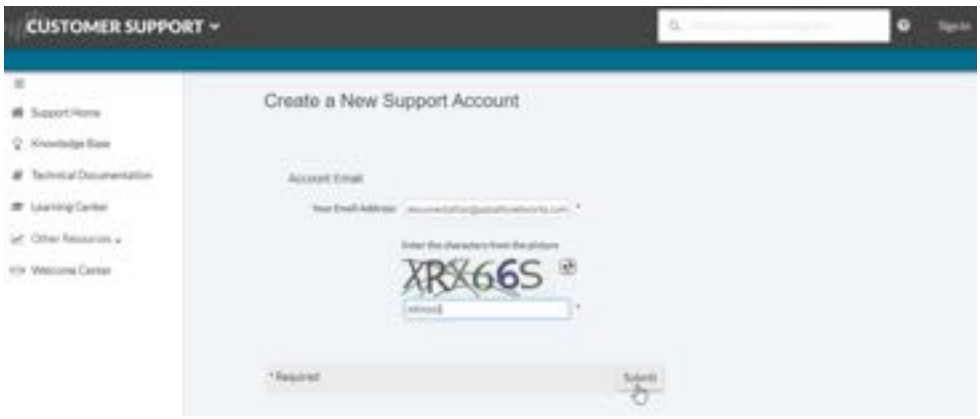
If you do not already have an active Palo Alto Networks support account, then you need to register your firewall when you create your new support account.

STEP 1 | Go to the Palo Alto Networks Customer Support Portal.

STEP 2 | Click **Activate My Account**.



STEP 3 | Enter **Your Email Address**, **Enter the characters from the picture**, and then **Submit** your entries.



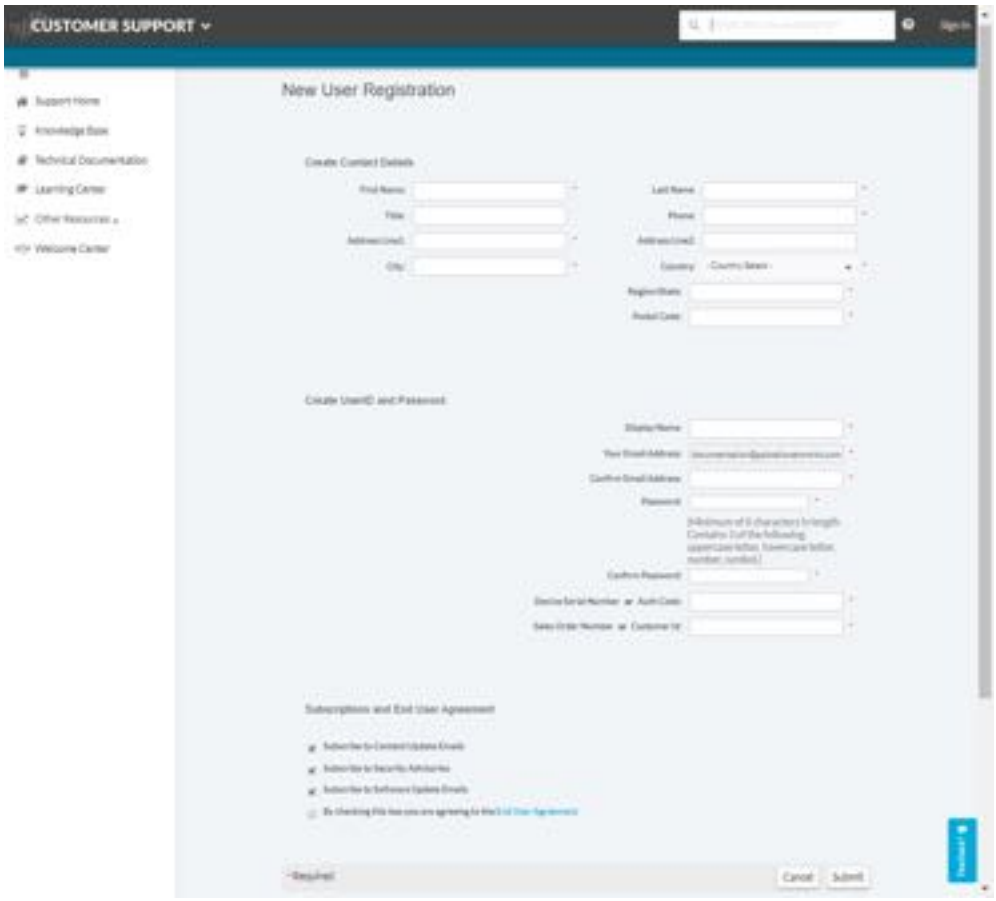
STEP 4 | Select **Register device using Serial Number or Authorization Code** and **Submit**.

© 2018 Palo Alto Networks, Inc.



STEP 5 | Complete the registration form.

1. Enter the contact details for the person in your organization who will own this account. Required fields are indicated by red asterisks.
2. Create a UserID and Password for the account. Required fields are indicated by a red asterisks.
3. Enter the Panorama **Serial Number** or **Auth Code**.
4. Enter your **Sales Order Number** or **Customer Id**.
5. To ensure that you are always alerted to the latest updates and security advisories, **Subscribe to Content Update Emails**, **Subscribe to Security Advisories**, and **Subscribe to Software Update Emails**.
6. Select the check box to agree to the End User Agreement and **Submit**.



© 2018 Palo Alto Networks, Inc.

Register a Firewall

If you already have an active Palo Alto Networks Customer Support account, perform the following tasks to register your firewall.

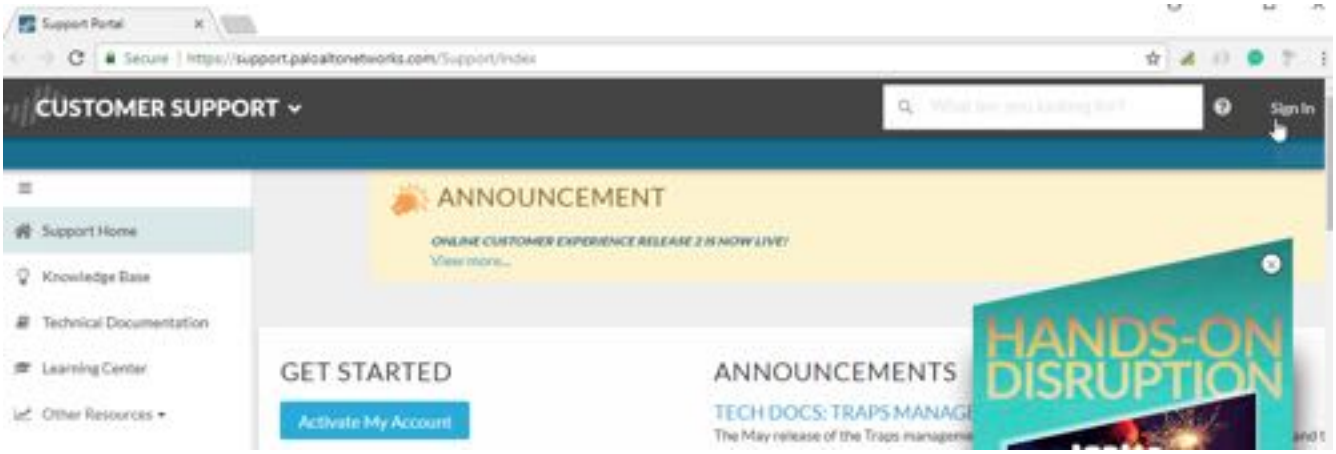
STEP 1 | Log in to the firewall web interface.

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

STEP 2 | Locate your serial number and copy it to the clipboard.

On the **Dashboard**, locate your **Serial Number** in the General Information section of the screen.

STEP 3 | Go to the Palo Alto Networks Customer Support portal and, if not already logged in, **Sign In** now.



STEP 4 | Register the firewall.

1. Click **Register a Device** in the Quick Actions section.



2. On the **Devices** page, click **Register New Device**.



3. Select **Register device using Serial Number or Authorization Code**, and then **Submit**.

© 2018 Palo Alto Networks, Inc.



4. Enter the firewall **Serial Number** (you can copy and paste it from the firewall Dashboard).
5. (Optional) Enter the **Device Name** and **Device Tag**.
6. (Optional) If the device will not have a connection to the internet, select the **Device will be used Offline** check box and then, from the drop-down, select the **OS Release** you plan to use.
7. Provide information about where you plan to deploy the firewall including the **City**, **Postal Code**, and **Country**.
8. Read the End User License Agreement ("EULA") and the Support Agreement and then **Agree and Submit**.



An entry for the firewall you just registered displays under **Devices**.

© 2018 Palo Alto Networks, Inc.

Activate Licenses and Subscriptions

Before you can start using your firewall to secure the traffic on your network, you must activate the licenses for each of the services you purchased. Available licenses and subscriptions include the following:

- Threat Prevention—Provides antivirus, anti-spyware, and vulnerability protection.
- Decryption Mirroring—Provides the ability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as NetWitness or Solera—for archiving and analysis.
- URL Filtering—Provides the ability to create security policy that allows or blocks access to the web based on dynamic URL categories. You must purchase and install a subscription for one of the supported URL filtering databases: PAN-DB or BrightCloud. With PAN-DB, you can set up access to the PAN-DB public cloud or to the PAN-DB private cloud. For more information about URL filtering, see Control Access to Web Content.
- Virtual Systems—This license is required to enable support for multiple virtual systems on PA-3000 Series firewalls. In addition, you must purchase a Virtual Systems license if you want to increase the number of virtual systems beyond the base number provided by default on PA-4000 Series, PA-5000 Series, PA-5200 Series, and PA-7000 Series firewalls (the base number varies by platform). The PA-800 Series, PA-500, PA-200, PA-220, and VM-Series firewalls do not support virtual systems.
- WildFire—Although basic WildFire support is included as part of the Threat Prevention license, the WildFire subscription service provides enhanced services for organizations that require immediate coverage for threats, frequent WildFire signature updates, advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet), as well as the ability to upload files using the WildFire API. A WildFire subscription is also required if your firewalls will be forwarding files to an on-premise WF-500 appliance.
- GlobalProtect—Provides mobility solutions and/or large-scale VPN capabilities. By default, you can deploy GlobalProtect portals and gateways (without HIP checks) without a license. If you want to use 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.
- AutoFocus—Provides a graphical analysis of firewall traffic logs and identifies potential risks to your network using threat intelligence from the AutoFocus portal. With an active license, you can also open an AutoFocus search based on logs recorded on the firewall.

STEP 1 | Locate the activation codes for the licenses you purchased.

When you purchased your subscriptions you should have received an email from Palo Alto Networks customer service listing the activation code associated with each subscription. If you cannot locate this email, contact Customer Support to obtain your activation codes before you proceed.

STEP 2 | Activate your Support license.

You will not be able to update your PAN-OS software if you do not have a valid Support license.

1. Log in to the web interface and then select **Device** > **Support**.
2. Click **Activate support using authorization code**.
3. Enter your **Authorization Code** and then click **OK**.

STEP 3 | Activate each license you purchased.

Select **Device** > **Licenses** and then activate your licenses and subscriptions in one of the following ways:

© 2018 Palo Alto Networks, Inc.

- **Retrieve license keys from license server**—Use this option if you activated your license on the Customer Support portal.
- **Activate feature using authorization code**—Use this option to enable purchased subscriptions using an authorization code for licenses that have not been previously activated on the support portal. When prompted, enter the **Authorization Code** and then click **OK**.
- **Manually upload license key**—Use this option if your firewall does not have connectivity to the Palo Alto Networks Customer Support Portal. In this case, you must download a license key file from the support site on an Internet connected computer and then upload to the firewall.

STEP 4 | Verify that the license was successfully activated

On the **Device** > **Licenses** page, verify that the license was successfully activated. For example, after activating the WildFire license, you should see that the license is valid:



STEP 5 | (WildFire subscriptions only) Perform a commit to complete WildFire subscription activation.

After activating a WildFire subscription, a commit is required for the firewall to begin forwarding advanced file types. You should either:

- Commit any pending changes.
- Check that the WildFire Analysis profile rules include the advanced file types that are now supported with the WildFire subscription. If no change to any of the rules is required, make a minor edit to a rule description and perform a commit.

© 2018 Palo Alto Networks, Inc.

Install Content and Software Updates

To ensure that you are always protected from the latest threats (including those that have not yet been discovered), you must ensure that you keep your firewalls up-to-date with the latest content and software updates published by Palo Alto Networks.

The following content updates are available, depending on which subscriptions you have:

- **Antivirus**—Includes new and updated antivirus signatures, including WildFire signatures and automatically-generated command-and-control (C2) signatures. WildFire signatures detect malware first seen by firewalls from around the world. Automatically-generated C2 detect certain patterns in C2 traffic (instead of the C2 server sending malicious commands to a compromised system); these signatures enable the firewall to detect C2 activity even when the C2 host is unknown or changes rapidly. You must have a Threat Prevention subscription to get these updates. New antivirus signatures are published daily.
- **Applications**—Includes new and updated application signatures. This update does not require any additional subscriptions, but it does require a valid maintenance/support contract. New application updates are published weekly. To ensure that Application updates do not impact your existing policy, review the policy impact of new application updates (see Manage New App-IDs Introduced in Content Releases) and be sure to follow the Best Practices for Application and Threat Content Updates.
- **Applications and Threats**—Includes new and updated application and threat signatures, including those that detect spyware and vulnerabilities. This update is available if you have a Threat Prevention subscription (and you get it instead of the Applications update). New Applications and Threats updates are published weekly, and the firewall can retrieve the latest update within 30 minutes of availability. To ensure that Application updates do not impact your existing policy, review the policy impact of new application updates (see Manage New App-IDs Introduced in Content Releases) and be sure to follow the Best Practices for Application and Threat Content Updates.
- **GlobalProtect Data File**—Contains the vendor-specific information for defining and evaluating host information profile (HIP) data returned by GlobalProtect agents. You must have a GlobalProtect license (subscription) and create an update schedule in order to receive these updates.
- **GlobalProtect Clientless VPN**—Contains new and updated application signatures to enable Clientless VPN access to common web applications from the GlobalProtect portal. You must have a GlobalProtect license (subscription) and create an update schedule in order to receive these updates and enable Clientless VPN to function.
- **BrightCloud URL Filtering**—Provides updates to the BrightCloud URL Filtering database only. You must have a BrightCloud subscription to get these updates. New BrightCloud URL database updates are published daily. If you have a PAN-DB license, scheduled updates are not required as firewalls remain in-sync with the servers automatically.
- **WildFire**—Available with a WildFire subscription, this update provides near real-time malware and antivirus signatures created as a result of the analysis done by the WildFire cloud service. As a best practice, schedule the firewall to retrieve and WildFire updates every minute (this allows the firewall to get the latest signatures within a minute of availability). If you have a Threat Prevention subscription and not a WildFire subscription, you must wait 24 to 48 hours for the WildFire signatures to roll into the antivirus update.
- **WF-Private**—Provides malware signatures generated by an on-premise WildFire appliance.

Follow the steps to install content and software updates.

STEP 1 | Ensure that the firewall has access to the update server.

© 2018 Palo Alto Networks, Inc.

1. By default, the firewall accesses the **Update Server** at **`updates.paloaltonetworks.com`** so that the firewall receives content updates from the server to which it is closest in the Content Delivery Network Infrastructure for Dynamic Updates. If the firewall has restricted access to the Internet, set the update server address to use the hostname **`staticupdates.paloaltonetworks.com`** or the IP address **`199.167.52.15`** instead of dynamically selecting a server from the CDN infrastructure.
2. (Optional) Click **Verify Update Server Identity** for an extra level of validation to enable the firewall to check that the server's SSL certificate is signed by a trusted authority. This is enabled by default.
3. (Optional) If the firewall needs to use a proxy server to reach Palo Alto Networks update services, in the **Proxy Server** window, enter:
   - **Server**—IP address or host name of the proxy server.
   - **Port**—Port for the proxy server. Range: 1-65535.
   - **User**—Username to access the server.
   - **Password**—Password for the user to access the proxy server. Re-enter the password at **Confirm Password**.

STEP 2 | Check for the latest content updates.

Select **Device** > **Dynamic Updates** and click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available:

- **Download**—Indicates that a new update file is available. Click the link to begin downloading the file directly to the firewall. After successful download, the link in the **Action** column changes from **Download** to **Install**.



You cannot download the antivirus update until you have installed the Application and Threats update.

- **Upgrade**—Indicates that a new version of the BrightCloud database is available. Click the link to begin the download and installation of the database. The database upgrade begins in the background; when completed a check mark displays in the **Currently Installed** column.

If you are using PAN-DB as your URL filtering database you will not see an upgrade link because the PAN-DB database on the firewall automatically synchronizes with the PAN-DB cloud.



To check the status of an action, click Tasks (on the lower right-hand corner of the window).

- **Revert**—Indicates that a previously installed version of the content or software version is available. You can choose to revert to the previously installed version.

STEP 3 | Install the content updates.

Installation can take up to 20 minutes on a PA-200 or PA-500 firewall and up to two minutes on a PA-5000 Series, PA-7000 Series, or VM-Series firewall.

© 2018 Palo Alto Networks, Inc.

Click the **Install** link in the **Action** column. When the installation completes, a check mark displays in the **Currently Installed** column.



STEP 4 | Schedule each content update.

Repeat this step for each update you want to schedule.

Stagger the update schedules because the firewall can only download one update at a time. If you schedule the updates to download during the same time interval, only the first download will succeed.

1. Set the schedule of each update type by clicking the **None** link.

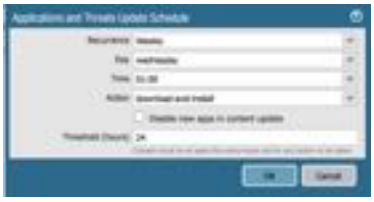


2. Specify how often you want the updates to occur by selecting a value from the **Recurrence** drop-down. The available values vary by content type (WildFire updates are available **Every Minute**, **Every 15 Minutes**, **Every 30 minutes**, or **Every Hour** whereas Applications and Threats updates can be scheduled for **Weekly**, **Daily**, **Hourly**, or **Every 30 Minutes** and Antivirus updates can be scheduled for **Hourly**, **Daily**, or **Weekly**).

   As new WildFire signatures are made available every five minutes, set the firewall to retrieve WildFire updates Every Minute to get the latest signatures within a minute of availability.

3. Specify the **Time** and (or, minutes past the hour in the case of WildFire), if applicable depending on the **Recurrence** value you selected, **Day** of the week that you want the updates to occur.
4. Specify whether you want the system to **Download Only** or, as a best practice, **Download And Install** the update.
5. Enter how long after a release to wait before performing a content update in the **Threshold (Hours)** field. In rare instances, errors in content updates may be found. For this reason, you may want to delay installing new updates until they have been released for a certain number of hours.

   If you have mission critical applications that must be 100% available, set the threshold for Applications or Applications and Threats updates to a minimum of 24 hours or more and follow the Best Practices for Application and Threat Content Updates*.*



6. Click **OK** to save the schedule settings.
7. Click **Commit** to save the settings to the running configuration.

STEP 5 | Update PAN-OS.

Always update content before updating PAN-OS. Every PAN-OS version has a minimum supported content release version*.*

1. Review the Release Notes.
2. Update the PAN-OS software.

© 2018 Palo Alto Networks, Inc.

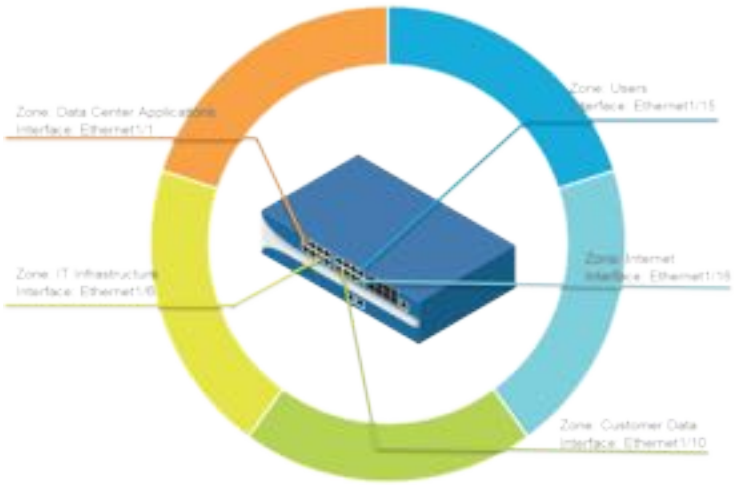
Segment Your Network Using Interfaces and Zones

Traffic must pass through the firewall in order for the firewall to manage and control it. Physically, traffic enters and exits the firewall through *interfaces*. The firewall determines how to act on a packet based on whether the packet matches a *Security policy rule*. At the most basic level, each Security policy rule must identify where the traffic came from and where it is going. On a Palo Alto Networks next-generation firewall, Security policy rules are applied between zones. A *zone* is a grouping of interfaces (physical or virtual) that represents a segment of your network that is connected to, and controlled by, the firewall. Because traffic can only flow between zones if there is a Security policy rule to allow it, this is your first line of defense. The more granular the zones you create, the greater control you have over access to sensitive applications and data and the more protection you have against malware moving laterally throughout your network. For example, you might want to segment access to the database servers that store your customer data into a zone called Customer Data. You can then define security policies that only permit certain users or groups of users to access the Customer Data zone, thereby preventing unauthorized internal or external access to the data stored in that segment.

- Network Segmentation for a Reduced Attack Surface on page 36
- Configure Interfaces and Zones on page 36

Network Segmentation for a Reduced Attack Surface

The following diagram shows a very basic example of network segmentation using zones. The more granular you make your zones (and the corresponding security policy rules that allows traffic between zones), the more you reduce the attack surface on your network. This is because traffic can flow freely within a zone (intra-zone traffic), but traffic cannot flow between zones (inter-zone traffic) until you define a Security policy rule that allows it. Additionally, an interface cannot process traffic until you have assigned it to a zone. Therefore, by segmenting your network into granular zones you have more control over access to sensitive applications or data and you can prevent malicious traffic from establishing a communication channel within your network, thereby reducing the likelihood of a successful attack on your network.



Configure Interfaces and Zones

After you identify how you want to segment your network and the zones you will need to create to achieve the segmentation (as well as the interfaces to map to each zone), you can begin configuring the interfaces

© 2018 Palo Alto Networks, Inc.

and zones on the firewall. Configure Interfaces on the firewall the to support the topology of each part of the network you are connecting to. The following workflow shows how to configure Layer 3 interfaces and assign them to zones. For details on integrating the firewall using a different type of interface deployments (for example as Virtual Wire Interfaces or as Layer 2 Interfaces), see Networking.

The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and virtual router). If you do not plan to use the default virtual wire, you must manually delete the configuration and commit the change before proceeding to prevent it from interfering with other settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in Set Up Network Access for External Services.

STEP 1 | Configure a default route to your Internet router.

1. Select **Network** > **Virtual Router** and then select the **default** link to open the Virtual Router dialog.
2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and enter the route in the **Destination** field (for example, 0.0.0.0/0).
3. Select the **IP Address** radio button in the **Next Hop** field and then enter the IP address and netmask for your Internet gateway (for example, 203.0.113.1).
4. Click **OK** twice to save the virtual router configuration.

STEP 2 | Configure the external interface (the interface that connects to the Internet).

1. Select **Network** > **Interfaces** and then select the interface you want to configure. In this example, we are configuring Ethernet1/16 as the external interface.
2. Select the **Interface Type**. Although your choice here depends on interface topology, this example shows the steps for **Layer3**.
3. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Internet, and then click **OK**.
4. In the **Virtual Router** drop-down, select **default**.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.113.23/24.
6. To enable you to ping the interface, select **Advanced** > **Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. Enter a **Name** for the profile, select **Ping** and then click **OK**.
7. To save the interface configuration, click **OK**.

STEP 3 | Configure the interface that connects to your internal network.

In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT.

1. Select **Network** > **Interfaces** and select the interface you want to configure. In this example, we are configuring Ethernet1/15 as the internal interface our users connect to.
2. Select **Layer3** as the **Interface Type**.
3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Users, and then click **OK**.
4. Select the same Virtual Router you used previously, default in this example.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24.
6. To enable you to ping the interface, select the management profile that you just created.
7. To save the interface configuration, click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 4 | Configure the interface that connects to your data center applications.

Although this basic security policy example configuration depicts using a single zone for all of your data center applications, as a best practice you would want to define more granular zones to prevent unauthorized access to sensitive applications or data and eliminate the possibility of malware moving laterally within your data center.

1. Select the interface you want to configure.
2. Select **Layer3** from the **Interface Type** drop-down. In this example, we are configuring Ethernet1/1 as the interface that provides access to your data center applications.
3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. In the Zone dialog, define a **Name** for new zone, for example Data Center Applications, and then click **OK**.
4. Select the same Virtual Router you used previously, default in this example.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24.
6. To enable you to ping the interface, select the management profile that you created.
7. To save the interface configuration, click **OK**.

STEP 5 | (Optional) Create tags for each zone.

Tags allow you to visually scan policy rules.

1. Select **Objects** > **Tags** and **Add**.
2. Select a zone **Name**.
3. Select a tag **Color** and click **OK**.



STEP 6 | Save the interface configuration.

Click **Commit**.

STEP 7 | Cable the firewall.

Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment.

STEP 8 | Verify that the interfaces are active.

Select **Dashboard** and verify that the interfaces you configured show as green in the Interfaces widget.



© 2018 Palo Alto Networks, Inc.

Set Up a Basic Security Policy

Now that you have defined some zones and attached them to interfaces, you are ready to begin creating your Security Policy. The firewall will not allow any traffic to flow from one zone to another unless there is a Security policy rule to allow 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, user, and the service. The firewall evaluates incoming traffic against the security policy rulebase from left to right and from top to bottom and then takes the action specified in the first security rule that matches (for example, whether to allow, deny, or drop the packet). This means that you must order the rules in your security policy rulebase so that more specific rules are at the top of the rulebase and more general rules are at the bottom to ensure that the firewall is enforcing policy as expected.

Even though a security policy rule allows a packet, this does not mean that the traffic is free of threats. To enable the firewall to scan the traffic that it allows based on a security policy rule, you must also attach Security Profiles—including URL Filtering, Antivirus, Anti-Spyware, File Blocking, and WildFire Analysis—to each rule (note that the profiles you can use depend on what subscriptions you have purchased). When creating your basic security policy, use the predefined security profiles to ensure that the traffic you allow into your network is being scanned for threats. You can customize these profiles later as needed for your environment.

Use following workflow set up a very basic security policy that enables access to the network infrastructure, to data center applications, and to the Internet. This will enable you to get the firewall up and running so that you can verify that you have successfully configured the firewall. This policy is not comprehensive enough to protect your network. After you verify that you have successfully configured the firewall and integrated it into your network, proceed with creating a best practice security policy that will safely enable application access while protecting your network from attack.

STEP 1 | (Optional) Delete the default security policy rule.

By default, the firewall includes a security rule named *rule1* that allows all traffic from Trust zone to Untrust zone. You can either delete the rule or modify the rule to reflect your zone naming conventions.

STEP 2 | Allow access to your network infrastructure resources.

1. Select **Policies** > **Security** and click **Add**.
2. Enter a descriptive **Name** for the rule in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Users.
4. In the **Destination** tab, set the **Destination Zone** to IT Infrastructure.



As a best practice, consider using address objects in the Destination Address field to enable access to specific servers or groups of servers only, particularly for services such as DNS and SMTP that are commonly exploited. By restricting users to specific destination server addresses you can prevent data exfiltration and command and control traffic from establishing communication through techniques such as DNS tunneling.

5. In the **Applications** tab, **Add** the applications that correspond to the network services you want to safely enable. For example, select **dns**, **ntp**, **ocsp**, **ping**, **smtp**.
6. In the **Service/URL Category** tab, keep the **Service** set to **application-default**.
7. In the **Actions** tab, set the **Action Setting** to **Allow**.
8. Set **Profile Type** to **Profiles** and select the following security profiles to attach to the policy rule:
   - For **Antivirus** select **default**

© 2018 Palo Alto Networks, Inc.

- For **Vulnerability Protection** select **strict**
- For **Anti-Spyware** select **strict**
- For **URL Filtering** select **default**
- For **File Blocking** select **basic file blocking**
- For **WildFire Analysis** select **default**

9. Verify that **Log at Session End** is enabled. Only traffic that matches a security policy rule will be logged.
10. Click **OK**.



STEP 3 | Enable access to general Internet applications.

This is a temporary rule that allows you to gather information about the traffic on your network. After you have more insight into what applications your users need access to, you can make informed decisions about what applications to allow and create more granular application-based rules for each user group.

1. Select **Policies** > **Security** and click **Add**.
2. Enter a descriptive **Name** for the rule in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Users.
4. In the **Destination** tab, set the **Destination Zone** to Internet.
5. In the **Applications** tab, **Add** an **Application Filter** and enter a Name. To safely enable access to legitimate web-based applications, set the Category in the application filter to **general-internet** and then click **OK**. To enable access to encrypted sites, **Add** the **ssl** application.
6. In the **Service/URL Category** tab, keep the **Service** set to **application-default**.
7. In the **Actions** tab, set the **Action Setting** to **Allow**.
8. Set **Profile Type** to **Profiles** and select the following security profiles to attach to the policy rule:
   - For **Antivirus** select **default**
   - For **Vulnerability Protection** select **strict**
   - For **Anti-Spyware** select **strict**
   - For **URL Filtering** select **default**
   - For **File Blocking** select **strict file blocking**
   - For **WildFire Analysis** select **default**
9. Verify that **Log at Session End** is enabled. Only traffic that matches a security rule will be logged.
10. Click **OK**.



STEP 4 | Enable access to data center applications.

1. Select **Policies** > **Security** and click **Add**.
2. Enter a descriptive **Name** for the rule in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Users.
4. In the **Destination** tab, set the **Destination Zone** to Data Center Applications.
5. In the **Applications** tab, **Add** the applications that correspond to the network services you want to safely enable. For example, select **activesync**, **imap**, **kerberos**, **ldap**, **ms-exchange**, and **ms-lync**.

© 2018 Palo Alto Networks, Inc.

6. In the **Service/URL Category** tab, keep the **Service** set to **application-default**.
7. In the **Actions** tab, set the **Action Setting** to **Allow**.
8. Set **Profile Type** to **Profiles** and select the following security profiles to attach to the policy rule:
   - For **Antivirus** select **default**
   - For **Vulnerability Protection** select **strict**
   - For **Anti-Spyware** select **strict**
   - For **URL Filtering** select **default**
   - For **File Blocking** select **basic file blocking**
   - For **WildFire Analysis** select **default**
9. Verify that **Log at Session End** is enabled. Only traffic that matches a security rule will be logged.
10. Click **OK**.



STEP 5 | Save your policy rules to the running configuration on the firewall.

Click **Commit**.

STEP 6 | To verify that you have set up your basic policies effectively, test whether your security policy rules are being evaluated and determine which security policy rule applies to a traffic flow.

To verify the policy rule that matches a flow, use the following CLI command:

```
test security-policy-match source <IP_address> destination <IP_address>
 destination port <port_number> application <application_name> protocol
 <protocol_number>
```

The output displays the best rule that matches the source and destination IP address specified in the CLI command.

For example, to verify the policy rule that will be applied for a client in the user zone with the IP address 10.35.14.150 when it sends a DNS query to the DNS server in the data center:

```
test security-policy-match source 10.35.14.150 destination 10.43.2.2
 application dns protocol 53

"Network Infrastructure" {
        from Users;
        source any;
        source-region none;
        to Data_Center;
        destination any;
        destination-region none;
        user any;
        category any;
        application/service  dns/any/any/any;
        action allow;
        icmp-unreachable: no
        terminal yes;
}
```


© 2018 Palo Alto Networks, Inc.

Assess Network Traffic

Now that you have a basic security policy, you can review the statistics and data in the Application Command Center (ACC), traffic logs, and the threat logs to observe trends on your network. Use this information to identify where you need to create more granular security policy rules.

- Use the Application Command Center and Use the Automated Correlation Engine.

  In the ACC, review the most used applications and the high-risk applications on your network. The ACC graphically summarizes the log information to highlight the applications traversing the network, who is using them (with User-ID enabled), and the potential security impact of the content to help you identify what is happening on the network in real time. You can then use this information to create appropriate security policy rules that block unwanted applications, while allowing and enabling applications in a secure manner.

  The Compromised Hosts widget in **ACC** > **Threat Activity** displays potentially compromised hosts on your network and the logs and match evidence that corroborates the events.

- Determine what updates/modifications are required for your network security policy rules and implement the changes.

  For example:

  - Evaluate whether to allow web content based on schedule, users, or groups.
  - Allow or control certain applications or functions within an application.
  - Decrypt and inspect content.
  - Allow but scan for threats and exploits.

  For information on refining your security policies and for attaching custom security profiles, see how to Create a Security Policy Rule and Security Profiles.

- View Logs.

  Specifically, view the traffic and threat logs (**Monitor** > **Logs**).

  Traffic logs are dependent on how your security policies are defined and set up to log traffic. The Application Usage widget in the ACC, however, records applications and statistics regardless of policy configuration; it shows all traffic that is allowed on your network, therefore it includes the inter-zone traffic that is allowed by policy and the same zone traffic that is allowed implicitly.

- Configure Log Storage Quotas and Expiration Periods.

  Review the AutoFocus intelligence summary for artifacts in your logs. An *artifact* is an item, property, activity, or behavior associated with logged events on the firewall. The intelligence summary reveals the number of sessions and samples in which WildFire detected the artifact. Use WildFire verdict information (benign, grayware, malware) and AutoFocus matching tags to look for potential risks in your network.

  AutoFocus tags created by Unit 42, the Palo Alto Networks threat intelligence team, call attention to advanced, targeted campaigns and threats in your network.

  From the AutoFocus intelligence summary, you can start an AutoFocus search for artifacts and assess their pervasiveness within global, industry, and network contexts.

© 2018 Palo Alto Networks, Inc.

- Monitor Web Activity of Network Users.

  Review the URL filtering logs to scan through alerts, denied categories/URLs. URL logs are generated when a traffic matches a security rule that has a URL filtering profile attached with an action of alert, continue, override or block.

© 2018 Palo Alto Networks, Inc.

Enable Basic WildFire Forwarding

WildFire is a cloud-based virtual environment that analyzes and executes unknown samples (files and email links) and determines the samples to be malicious, phishing, grayware, or benign. With WildFire enabled, a Palo Alto Networks firewall can forward unknown samples to WildFire for analysis. For newly-discovered malware, WildFire generates a signature to detect the malware and distributes it to all firewalls with active WildFire subscription within minutes. This enables all Palo Alto next-generation firewalls worldwide to detect and prevent malware found by a single firewall. When you enable WildFire forwarding, the firewall also forwards files that were blocked by Antivirus signatures, in addition to unknown samples. Malware signatures often match multiple variants of the same malware family, and as such, block new malware variants that the firewall has never seen before. The Palo Alto Networks threat research team uses the threat intelligence gathered from malware variants to block malicious IP addresses, domains, and URLs.

A basic WildFire service is included as part of the Palo Alto Networks next-generation firewall and does not require a WildFire subscription. With the basic WildFire service, you can enable the firewall to forward portable executable (PE) files. Additionally, if you do not have a WildFire subscription, but you do have a Threat Prevention subscription, you can receive signatures for malware WildFire identifies every 24- 48 hours (as part of the Antivirus updates).

Beyond the basic WildFire service, a WildFire subscription is required for the firewall to:

- Get the latest WildFire signatures every five minutes.
- Forward advanced file types and email links for analysis.
- Use the WildFire API.
- Use a WildFire appliance to host a WildFire private cloud or a WildFire hybrid cloud.

If you have a WildFire subscription, go ahead and get started with WildFire to get the most out of your subscription. Otherwise, take the following steps to enable basic WildFire forwarding:

STEP 1 | Confirm that your firewall is registered and that you have a valid support account as well as any subscriptions you require.

1. Log in to the Palo Alto Networks Customer Support Portal(CSP). **My Devices**.
2. On the left-hand navigation pane, select **Assets** > **Devices** and verify that the firewall is listed. If it's not listed, select **Register New Device** and Register the Firewall.
3. (Optional) If you have a Threat Prevention subscription, be sure to Activate Licenses and Subscriptions.

STEP 2 | Configure WildFire forwarding settings.

1. Select **Device** > **Setup** > **WildFire** and edit the General Settings.
2. Set the **WildFire Public Cloud** field to forward files to the WildFire global cloud at: **`wildfire.paloaltonetworks.com`**.

   You can also forward files to a regional cloud or a private cloud based on your location and your organizational requirements.

3. Review the **File Size Limits** for PEs the firewall forwards for WildFire analysis. set the **Size Limit** for PEs that the firewall can forward to the maximum available limit of 10 MB.

   As a WildFire best practice, set the Size Limit for PEs to the maximum available limit of 10 MB.

4. Click **OK** to save your changes.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Enable the firewall to forward PEs for analysis.

1. Select **Objects** > **Security Profiles** > **WildFire Analysis** and **Add** a new profile rule.
2. **Name** the new profile rule.
3. **Add** a forwarding rule and enter a **Name** for it.
4. In the **File Types** column, add **pe** files to the forwarding rule.
5. In the **Analysis** column, select **public-cloud** to forward PEs to the WildFire public cloud.
6. Click **OK**.

STEP 4 | Apply the new WildFire Analysis profile to traffic that the firewall allows.

1. Select **Policies** > **Security** and either select an existing policy rule or create a new policy rule as described in Set Up a Basic Security Policy.
2. Select **Actions** and in the Profile Settings section, set the **Profile Type** to **Profiles**.
3. Select the **WildFire Analysis** profile you just created to apply that profile rule to all traffic this policy rule allows.
4. Click **OK**.

STEP 5 | Enable the firewall to forward decrypted SSL traffic for WildFire analysis.

STEP 6 | Review and implement WildFire best practices to ensure that you are getting the most of WildFire detection and prevention capabilities.

STEP 7 | **Commit** your configuration updates.

STEP 8 | Verify that the firewall is forwarding PE files to the WildFire public cloud.

Select **Monitor** > **Logs** > **WildFire Submissions** to view log entries for PEs the firewall successfully submitted for WildFire analysis. The Verdict column displays whether WildFire found the PE to be malicious, grayware, or benign. (WildFire only assigns the phishing verdict to email links). The Action column indicates whether the firewall allowed or blocked the sample. The Severity column indicates how much of a threat a sample poses to an organization using the following values: critical, high, medium, low, information.

STEP 9 | (Threat Prevention subscription only) If you have a Threat Prevention subscription, but do not have a WildFire subscription, you can still receive WildFire signature updates every 24- 48 hours.

1. Select **Device** > **Dynamic Updates**.
2. Check that the firewall is scheduled to download, and install Antivirus updates.

© 2018 Palo Alto Networks, Inc.

Control Access to Web Content

URL Filtering provides visibility and control over web traffic on your network. With URL filtering enabled, the firewall can categorize web traffic into one or more URL categories. You can then create policies that specify whether to allow, block, or log (alert) traffic based on the category to which it belongs. Together with User-ID, you can also use URL Filtering to Prevent Credential Phishing based on URL category.

The following workflow shows how to enable PAN-DB for URL filtering, create security profiles, and attach them to Security policy rules to enforce a basic URL filtering policy.

STEP 1 | Confirm that you have a URL Filtering license.

1. Obtain and install a URL Filtering license. See Activate Licenses and Subscriptions for details.
2. Select **Device** > **Licenses** and verify that the URL Filtering license is valid.



STEP 2 | Download the seed database and activate the license.

1. To download the seed database, click **Download** next to **Download Status** in the PAN-DB URL Filtering section of the Licenses page.
2. Choose a region (APAC, Europe, Japan, Latin-America, North-America, or Russia) and then click **OK** to start the download.
3. After the download completes, click **Activate**. The Active field now shows that PAN-DB is now active.



STEP 3 | Configure URL Filtering.

Configure a best practice URL Filtering profile to ensure protection against URLs that have been observed hosting malware or exploitive content.

Select **Objects** > **Security Profiles** > **URL Filtering** and **Add** or modify a URL Filtering profile.

- Select **Categories** to allow, alert, continue, or block access to. If you are not sure what sites or categories you want to control access to, consider setting the categories (except for those blocked by default) to alert. You can then use the visibility tools on the firewall, such as the ACC and App Scope, to determine which web categories to restrict to specific groups or to block entirely. See URL Filtering Profile Actions for details on the site access settings you can enforce for each URL category.
- Select **Categories** to Prevent Credential Phishing based on URL category.
- Select **Overrides** to Allow Password Access to Certain Sites.
- Enable Safe Search Enforcement to ensure that user search results are based on search engine safe search settings.

STEP 4 | Attach the URL filtering profile to a Security policy rule.

1. Select **Policies** > **Security**.

© 2018 Palo Alto Networks, Inc.

2. Select a Security policy rule that allows web access to edit it and select the **Actions** tab.
3. In the **Profile Settings** list, select the **URL Filtering** profile you just created. (If you don't see drop-downs for selecting profiles, set the **Profile Type** to **Profiles**.)
4. Click **OK** to save the profile.

STEP 5 | Enable response pages in the management profile for each interface on which you are filtering web traffic.

1. Select **Network** > **Network Profiles** > **Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.
2. Select **Response Pages**, as well as any other management services required on the interface.
3. Click **OK** to save the interface management profile.
4. Select **Network** > **Interfaces** and select the interface to which to attach the profile.
5. On the **Advanced** > **Other Info** tab, select the interface management profile you just created.
6. Click **OK** to save the interface settings.

STEP 6 | Commit your changes.

Commit the configuration.

STEP 7 | Test the URL filtering configuration.

From an endpoint in a trusted zone, attempt to access sites in various categories and make sure you see the expected result based on the corresponding Site Access setting you selected:

- If you set Site Access to **alert** for the category, check the URL Filtering log to make sure you see a log entry for the request.
- If you set Site Access to **continue** for the category, verify that the URL Filtering Continue and Override Page response page displays. **Continue** to the site.
- If you set Site Access to **block** for the category, verify that the URL Filtering and Category Match Block Page response page displays:



© 2018 Palo Alto Networks, Inc.

Enable AutoFocus Threat Intelligence

With a valid AutoFocus subscription, you can compare the activity on your network with the latest threat data available on the AutoFocus portal. Connecting your firewall and AutoFocus unlocks the following features:

- Ability to view an AutoFocus intelligence summary for session artifacts recorded in the firewall logs.
- Ability to open an AutoFocus search for log artifacts from the firewall.

The AutoFocus intelligence summary reveals the prevalence of an artifact on your network and on a global scale. The WildFire verdicts and AutoFocus tags listed for the artifact indicate whether the artifact poses a security risk.

STEP 1 | Verify that the AutoFocus license is activated on the firewall.

1. Select **Device** > **Licenses** to verify that the AutoFocus Device License is installed and valid (check the expiration date).
2. If the firewall doesn't detect the license, see Activate Licenses and Subscriptions.

STEP 2 | Connect the firewall to AutoFocus.

1. Select **Device** > **Setup** > **Management** and edit the AutoFocus settings.
2. Enter the **AutoFocus URL**:

   `https://autofocus.paloaltonetworks.com:10443`
3. Use the **Query Timeout** field to set the duration of time for the firewall to attempt to query AutoFocus for threat intelligence data. If the AutoFocus portal does not respond before the end of the specified period, the firewall closes the connection.

   As a best practice, set the query timeout to the default value of 15 seconds. AutoFocus queries are optimized to complete within this duration.
4. Select **Enabled** to allow the firewall to connect to AutoFocus.
5. Click **OK**.
6. **Commit** your changes to retain the AutoFocus settings upon reboot.

STEP 3 | Connect AutoFocus to the firewall.

1. Log in to the AutoFocus portal: https://autofocus.paloaltonetworks.com
2. Select **Settings**.
3. **Add new** remote systems.
4. Enter a descriptive **Name** to identify the firewall.
5. Select **PanOS** as the System Type.
6. Enter the firewall IP **Address**.
7. Click **Save changes** to add the remote system.
8. Click **Save changes** again on the Settings page to ensure the firewall is successfully added.

STEP 4 | Test the connection between the firewall and AutoFocus.

1. On the firewall, select **Monitor** > **Logs** > **Traffic**.
2. Verify that you can Assess Firewall Artifacts with AutoFocus.

© 2018 Palo Alto Networks, Inc.

Best Practices for Completing the Firewall Deployment

Now that you have integrated the firewall into your network and enabled the basic security features, you can begin configuring more advanced features. Here are some things to consider next:

- Follow the Best Practices for Securing Administrative Access to make sure you are properly securing the management interfaces.
- Configure a best-practice security policy rulebase to safely enable applications and protect your network from attack. Go to the Best Practices page and select security policy best practice for your firewall deployment.
- Set up High Availability—High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration and session tables are synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up a two-firewall cluster provides redundancy and allows you to ensure business continuity.
- Enable User Identification (User-ID)—User-ID is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses.
- Enable Decryption—Palo Alto Networks firewalls provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted or tunneled traffic.
- Follow the Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions.
- Share Threat Intelligence with Palo Alto Networks—Permit the firewall to periodically collect and send information about applications, threats, and device health to Palo Alto Networks. Telemetry includes options to enable passive DNS monitoring and to allow experimental test signatures to run in the background with no impact to your security policy rules, firewall logs, or firewall performance. All Palo Alto Networks customers benefit from the intelligence gathered from telemetry, which Palo Alto Networks uses to improve the threat prevention capabilities of the firewall.

© 2018 Palo Alto Networks, Inc.

Best Practices for Securing Administrative Access

Protecting your network from cyberattacks begins with a secure firewall deployment. If the network you use to manage your sensitive IT devices—including your Palo Alto Networks next-gen firewalls and Panorama—is not secured properly, you will not be able to detect and defend against vulnerability exploits that could lead to infiltration and/or the loss of sensitive data. The ultimate goal when securing firewall access is to ensure that even if an attacker gains access to privileged credentials, you can still thwart their ability to get in and do damage. Follow these best practice guidelines to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.

- Isolate the management network.

  All Palo Alto Networks firewalls provide an out-of-band management port (MGT) that you can use to perform the firewall administration functions. Alternatively, you can choose to use the MGT port for initial configuration, and then configure a data port for management access to the firewall. Either way, because the management interface provides access to your security configuration, you must take the following precautions to safeguard access to this interface:

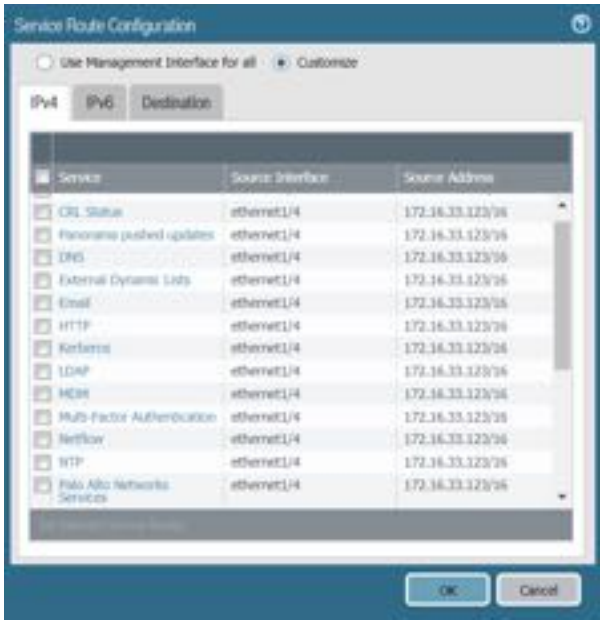
  Do not enable access to your management interface from the internet or from other untrusted zones inside your enterprise security boundary. This applies whether you use the dedicated management port (MGT) or you configured a data port as your management interface.

  - Isolate the management interface on a dedicated management VLAN.
  - Limit the source IP addresses allowed in to the management network to those of your dedicated management devices, such as a jump server or a bastion host.
  - Use a jump server or bastion host (with screen recording) to provide secure access from your corporate network in to your dedicated management network, and require that users authenticate and are authorized to access your management network.
  - If you don't have a bastion host, use Authentication Policy with multi-factor authentication (MFA) to require administrators to successfully authenticate before you allow them to continue to the firewall web interface login page or CLI login prompt. This prevents access to the management interface using stolen credentials or through vulnerability exploits.
  - Limit access to users in your security admin, network admin, or IT user groups, as appropriate for your organization.
  - If you must enable remote access to the management network, require access through a VPN tunnel using GlobalProtect. After administrators successfully establish a VPN tunnel into your VPN zone, they must still authenticate into the management network through your bastion host.
  - Do not use an interface management profile that allows HTTP, HTTPS, Telnet, or SSH on the interface where you have configured a GlobalProtect portal or gateway because this configuration exposes access to the management interface via the internet.
  - If you are using a template to deploy a VM-Series firewall that includes a field for restricting management access to a specific IP address, make sure to supply a CIDR block that corresponds to your dedicated management IP addresses or network. If necessary, modify the corresponding security group to add additional hosts or networks after the template launch. Do not make the allowed source network range larger than necessary and do not ever configure the allowed source as 0.0.0.0/0.

© 2018 Palo Alto Networks, Inc.

- Use service routes (**Device** > **Setup** > **Services** > **Service Route Configuration**) to enable the firewall to access services outside of the management network.

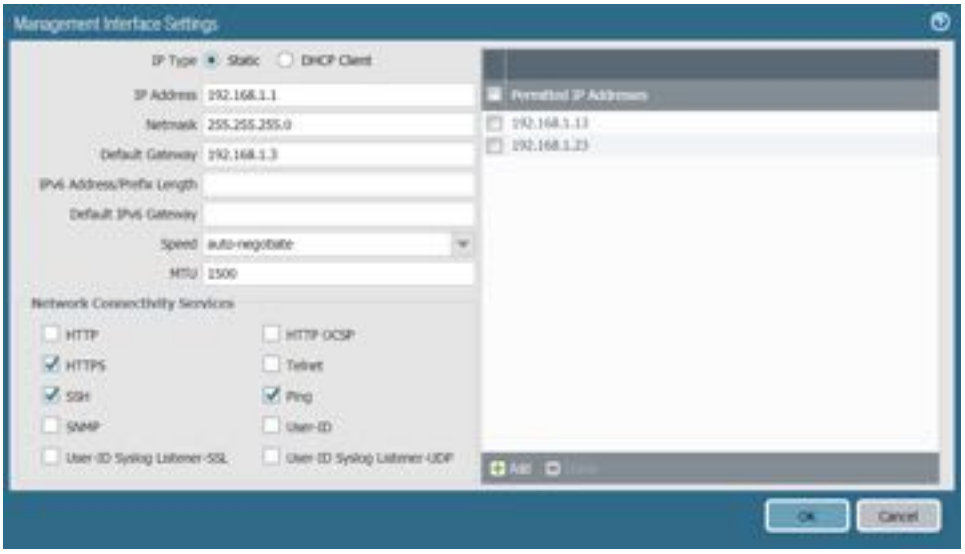
  By default, the firewall uses the management (MGT) port to access services—such as DNS servers, NTP servers, and authentication servers—that are on potentially untrusted networks, including services that require internet access, such as Palo Alto Networks Services and AutoFocus. Because your management interface—whether on the MGT port or a data port—must be isolated on the management network, you must use service routes to enable access to these services. When you configure a service route, the firewall instead uses the specified source interface and address to access the services you need. Specify the source IP address/interface for your service route on an interface that does not have management access (HTTPS or SSH) enabled.



- Restrict access to the management interface.
  - Restrict the IP addresses that are permitted to access the management interface.

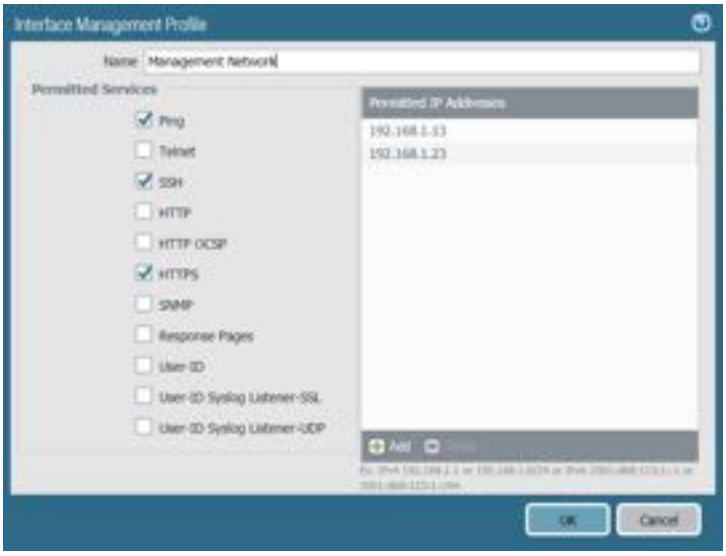
    Even if your firewall is on a dedicated management network that is only accessible by a device on the same VLAN or through a bastion host or VPN tunnel, you can secure the firewall further by restricting the source IP addresses that can access the management interface to those of your administrators. Limiting access to the management interface reduces the attack surface by helping to prevent access from unexpected IP addresses or subnets and prevents access using stolen credentials.
  - Restrict the services that are available on the management interface.
    - Do not allow access over Telnet and HTTP because these services use plaintext and are not as secure as the other services and could compromise administrator credentials. Instead, require administrators to access the firewall interfaces over SSH or HTTPS.
    - Enable ping if you want to be able to test connectivity to the interface, but do not enable any other services on the management interface.
  - The way you configure these settings depends on whether you are using the MGT port or a data port for access to the firewall management interfaces:
    - If you are using the MGT port as your management interface, select **Device** > **Setup** > **Interfaces** and select the **Management** interface to configure the settings to restrict who can access the management interface and what services the interface allows.

© 2018 Palo Alto Networks, Inc.



- If you are using a data port as your management interface, after you configure the interface, select **Network** > **Network Profiles** > **Interface Mgmt** and **Add** an interface management profile to restrict who can access the management interface and what services the interface allows.

  Do not attach an interface management profile that allows Telnet, SSH, HTTP, or HTTPS to an interface where you have configured a GlobalProtect portal or gateway because this will expose the management interface to the internet.



- Manage administrator access.
  - The firewall is preconfigured with a default administrative account (admin) that provides full read-write access (also known as superuser access) to the firewall. You must change the default admin account password (**Device** > **Administrators** > **admin**) immediately upon initial configuration.
  - Configure a Firewall Administrator Account for each individual who needs access to the administrative or reporting functions of the firewall. This allows you to better protect the firewall from unauthorized configuration (or modification) and to enable logging of the actions of each individual administrator.

© 2018 Palo Alto Networks, Inc.

- Assign each administrator account to an admin role profile that limits management privileges to only those functions the individual administrator needs.
- For any administrators with change privileges, require multi-factor authentication (MFA) using external authentication and authorization using RADIUS or SAML. See Configure Local or External Authentication for Firewall Administrators for details on how to configure external authentication with MFA.

  If you have a strong authentication infrastructure using smart cards, Configure Certificate-Based Administrator Authentication to the Web Interface *and* Configure SSH Key-Based Administrator Authentication to the CLI*.*

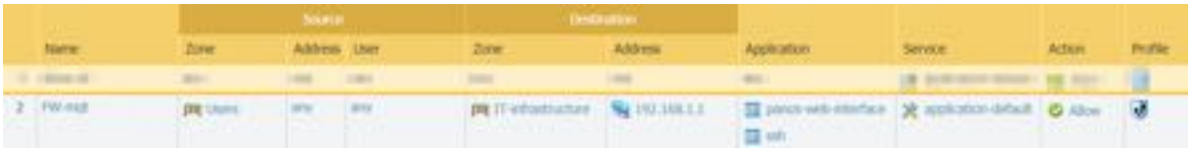
  If available, use privileged account management (PAM) and/or privileged identity management (PIM) solutions to secure administrator credentials externally.
- Configure a strict password policy, including requiring frequent password changes (**Device** > **Setup** > **Management** > **Minimum Password Complexity**).
- Monitor the System logs to identify abnormal account activity on any of your administrator accounts. For example, if the logs show excessive login attempts or repeated logins at certain times of day, this may indicate that an administrative account has been compromised. Also, educate all administrators about how to Use the Administrator Login Activity Indicators to Detect Account Misuse.

- Scan all traffic destined for the management interface for threats.

  Because security policy and decryption policy do not evaluate management plane traffic, you cannot directly scan the MGT port for threats. If you are using the MGT port as your management interface, consider routing traffic destined for the MGT port through a data port or through another firewall so that you can apply these important security checks to your management traffic.

  - Create security policy rules to allow access to the management interfaces of the firewall and Panorama (web interface or CLI). The way you define the policy depends on whether or not you are using a bastion host to enable access to the management network.
    - If you are not using a bastion host to isolate your management network, create a security policy rule to allow access from the Users zone to the IT Infrastructure zone. This security policy rule must be very granular and specify the source zone, source IP address (if available), and source user group of the user attempting to access the management interface, as well as the destination zone, IP address of the appliance (firewall or Panorama), and the App-ID to identify the specific management application (web interface or CLI) running on the application default port. For example, you would use the panos-web-interface App-ID to allow access to the web interface and the ssh App-ID to allow access to the CLI. You must also attach a Vulnerability Protection profile to the rule, as described in the next section.

      The following example rule allows access from the users zone directly to the IT Infrastructure zone, and restricts access to users in the IT-admins group who are attempting to access the management interface IP address to access the panos-web-interface application on the application-default port only:

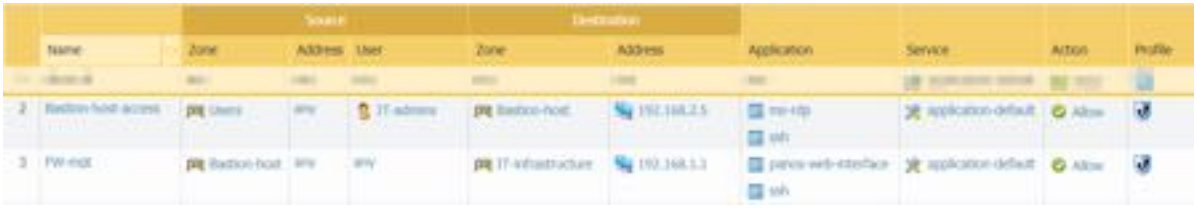
      

    - If you are using a bastion host to enable access to your management network, you need two security policy rules: one rule to allow access from user zone to the bastion host zone and a second rule to allow access from the bastion host zone to the IT infrastructure zone. Again, both of these security policy rules should be as granular as possible and include source zone, address

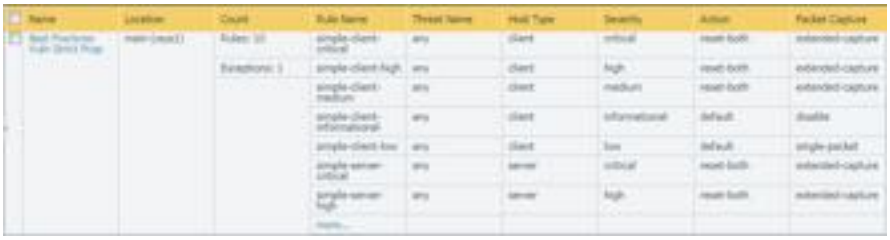
© 2018 Palo Alto Networks, Inc.

(if available), user and destination zone and address, and the App-ID. Keep in mind that if you are using a bastion host, the user IP address is usually the IP address of the bastion host so you cannot identify the User-ID for the administrator unless you are using the TerminalServicesagent on the bastion host to identify individual users. In this case, you must also attach a Vulnerability Protection profile to both rules, as described in the next section.

In the example rules that follow, the first rule allows access from the Users zone to the Bastion-host zone for users in the IT-admins group who are attempting to access the specified bastion server IP address over SSH and/or RDP. The second rule allows access to users from the Bastion-host zone to the IT-infrastructure zone attempting to access the panos-web-interface application on the default port on the firewall with the specified destination address.



- Attach a best practice Vulnerability Protection profile to the security policy rules that allow access in to your management network to protect against buffer overflows, illegal code execution, and other attempts to exploit client- and server-side vulnerabilities. To create a profile for the purpose of protecting your management interface, clone the strict profile and then enable extended packet capture to help you track down the source of any potential attacks.



- Configure SSL Inbound Inspection or Configure SSL Forward Proxy for traffic to or from the management interface to ensure that you can scan the traffic for threats. Attach a best practice decryption profile to the decryption policy rule to ensure that you are blocking vulnerable SSL/TLS versions such as TLS 1.0 and SSLv3, and to rejecting sessions using weak encryption algorithms such as RC4 and 3DES, and weak authentication algorithms such as MD5 and SHA1

- Replace the Certificate for Inbound Management Traffic.

  By default, the firewall ships with a default certificate that enables HTTPS access to the web interface over the management (MGT) interface or any other interface that supports HTTPS management traffic. To improve the security of inbound management traffic, replace the default certificate with a new certificate issued specifically for your organization. Use certificates signed by your enterprise CA so that users won't learn to ignore certificate warnings. In addition, in the SSL/TLS profile, set the **Min version** to **TLSv1.2**.

- Keep the firewall up-to-date with content and software updates to ensure that you are always protected by the latest security patches and threat updates.

  - To ensure that you are always alerted to the latest updates and security advisories, go to the Palo Alto Networks Support Portal, select **Edit Profile**, and **Subscribe to Content Update Emails**, **Subscribe to Security Advisories**, and **Subscribe to Software Update Emails**. Make sure you **Save Edits**.

© 2018 Palo Alto Networks, Inc.



- Follow the Best Practices for Application and Threat Content Updates when updating to the latest content release version.
- Before you upgrade PAN-OS, read the latest release notes.

© 2018 Palo Alto Networks, Inc.

Firewall Administration

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

> Management Interfaces
> Use the Web Interface
> Manage Configuration Backups
> Manage Firewall Administrators
> Reference: Web Interface Administrator Access
> Reference: Port Number Usage
> Reset the Firewall to Factory Default Settings
> Bootstrap the Firewall

© 2018 Palo Alto Networks, Inc.

Management Interfaces

You can use the following user interfaces to manage the Palo Alto Networks firewall:



Do not enable management access from the internet or from other untrusted zones inside your enterprise security boundary. Follow the Best Practices for Securing Administrative Access *to ensure that you are properly securing your firewall.*

- Use the Web Interface to perform configuration and monitoring tasks with relative ease. This graphical interface allows you to access the firewall using HTTPS (recommended) or HTTP and it is the best way to perform administrative tasks.
- Use the Command Line Interface (CLI) to perform a series of tasks by entering commands in rapid succession over SSH (recommended), Telnet, or the console port. The CLI is a no-frills interface that supports two command modes, operational and configure, each with a distinct hierarchy of commands and statements. When you become familiar with the nesting structure and syntax of the commands, the CLI provides quick response times and administrative efficiency.
- Use the XML API to streamline your operations and integrate with existing, internally developed applications and repositories. The XML API is a web service implemented using HTTP/HTTPS requests and responses.
- Use Panorama to perform web-based management, reporting, and log collection for multiple firewalls. The Panorama web interface resembles the firewall web interface but with additional functions for centralized management.

© 2018 Palo Alto Networks, Inc.

Use the Web Interface

The following topics describe how to use the firewall web interface. For detailed information about specific tabs and fields in the web interface, refer to the Web Interface Reference Guide.

- Launch the Web Interface
- Configure Banners, Message of the Day, and Logos
- Use the Administrator Login Activity Indicators to Detect Account Misuse
- Manage and Monitor Administrative Tasks
- Commit, Validate, and Preview Firewall Configuration Changes
- Use Global Find to Search the Firewall or Panorama Management Server
- Manage Locks for Restricting Configuration Changes

Launch the Web Interface

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

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

Perform the following tasks to launch the web interface.

STEP 1 | Launch an Internet browser and enter the IP address of the firewall in the URL field (https://<IP address>).

By default, the management (MGT) interface allows only HTTPS access to the web interface. To enable other protocols, select Device > Setup > Interfaces and edit the Management interface.

STEP 2 | Log in to the firewall according to the type of authentication used for your account. If logging in to the firewall for the first time, use the default value **admin** for your username and password.

- **SAML**—Click **Use Single Sign-On** (SSO). If the firewall performs authorization (role assignment) for administrators, enter your **Username** and **Continue**. If the SAML identity provider (IdP) performs authorization, **Continue** without entering a **Username**. In both cases, the firewall redirects you to the IdP, which prompts you to enter a username and password. After you authenticate to the IdP, the firewall web interface displays.
- **Any other type of authentication**—Enter your user **Name** and **Password**. Read the login banner and select **I Accept and Acknowledge the Statement Below** if the login page has the banner and check box. Then click **Login**.

STEP 3 | Read and **Close** the messages of the day.

Configure Banners, Message of the Day, and Logos

A *login banner* is optional text that you can add to the login page so that administrators will see information they must know before they log in. For example, you could add a message to notify users of restrictions on unauthorized use of the firewall.

© 2018 Palo Alto Networks, Inc.

You can add colored bands that highlight overlaid text across the top (*header banner*) and bottom (*footer banner*) of the web interface to ensure administrators see critical information, such as the classification level for firewall administration.

A *message of the day* dialog automatically displays after you log in. The dialog displays messages that Palo Alto Networks embeds to highlight important information associated with a software or content release. You can also add one custom message to ensure administrators see information, such as an impending system restart, that might affect their tasks.

You can replace the default logos that appear on the login page and in the header of the web interface with the logos of your organization.

STEP 1 | Configure the login banner.

1. Select **Device** > **Setup** > **Management** and edit the General Settings.
2. Enter the **Login Banner** (up to 3,200 characters).
3. (Optional) Select **Force Admins to Acknowledge Login Banner** to force administrators to select an **I Accept and Acknowledge the Statement Below** check box above the banner text to activate the **Login** button.
4. Click **OK**.

STEP 2 | Set the message of the day.

1. Select **Device** > **Setup** > **Management** and edit the Banners and Messages settings.
2. Enable the **Message of the Day**.
3. Enter the **Message of the Day** (up to 3,200 characters).

   After you enter the message and click OK, administrators who subsequently log in, and active administrators who refresh their browsers, see the new or updated message immediately; a commit isn't necessary. This enables you to inform other administrators of an impending commit that might affect their configuration changes. Based on the commit time that your message specifies, the administrators can then decide whether to complete, save, or undo their changes.
4. (Optional) Select **Allow Do Not Display Again** (default is disabled) to give administrators the option to suppress a message of the day after the first login session. Each administrator can suppress messages only for his or her own login sessions. In the message of the day dialog, each message will have its own suppression option.
5. (Optional) Enter a header **Title** for the message of the day dialog (default is `Message of the Day`).

STEP 3 | Configure the header and footer banners.

A bright background color and contrasting text color can increase the likelihood that administrators will notice and read a banner. You can also use colors that correspond to classification levels in your organization.

1. Enter the **Header Banner** (up to 3,200 characters).
2. (Optional) Clear **Same Banner Header and Footer** (enabled by default) to use different header and footer banners.
3. Enter the **Footer Banner** (up to 3,200 characters) if the header and footer banners differ.
4. Click **OK**.

STEP 4 | Replace the logos on the login page and in the header.

The maximum size for any logo image is 128KB. The supported file types are png, gif, and jpg. The firewall does not support image files that are interlaced or that contain alpha channels.

1. Select **Device** > **Setup** > **Operations** and click **Custom Logos** in the Miscellaneous section.

© 2018 Palo Alto Networks, Inc.

2. Perform the following steps for both the **Login Screen** logo and the **Main UI** (header) logo:
    1. Click upload .
    2. Select a logo image and click **Open**.

        You can preview the image to see how PAN-OS will crop it to fit by clicking the magnifying glass icon.
    3. Click **Close**.
3. **Commit** your changes.

STEP 5 | Verify that the banners, message of the day, and logos display as expected.

1. Log out to return to the login page, which displays the new logos you selected.
2. Enter your login credentials, review the banner, select **I Accept and Acknowledge the Statement Below** to enable the **Login** button, and then **Login**.

    A dialog displays the message of the day. Messages that Palo Alto Networks embedded display on separate pages in the same dialog. To navigate the pages, click the right or left arrows along the sides of the dialog or click a page selector at the bottom of the dialog.
3. (Optional) You can select **Do not show again** for the message you configured and for any messages that Palo Alto Networks embedded.
4. **Close** the message of the day dialog to access the web interface.

    Header and footer banners display in every web interface page with the text and colors that you configured. The new logo you selected for the web interface displays below the header banner.

Use the Administrator Login Activity Indicators to Detect Account Misuse

The last login time and failed login attempts indicators provide a visual way to detect misuse of your administrator account on a Palo Alto Networks firewall or Panorama management server. Use the last login information to determine if someone else logged in using your credentials and use the failed login attempts indicator to determine if your account is being targeted in a brute-force attack.

STEP 1 | View the login activity indicators to monitor recent activity on your account.

1. Log in to the web interface on your firewall or Panorama management server.
2. View the last login details located at the bottom left of the window and verify that the timestamp corresponds to your last login.



3. Look for a caution symbol to the right of the last login time information for failed login attempts.

    The failed login indicator appears if one or more failed login attempts occurred using your account since the last successful login.
    1. If you see the caution symbol, hover over it to display the number of failed login attempts.

© 2018 Palo Alto Networks, Inc.

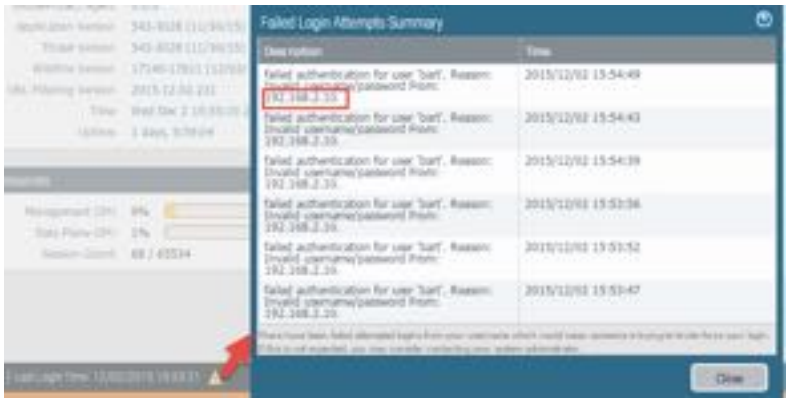


2. Click the caution symbol to view the failed login attempts summary. Details include the admin account name, the reason for the login failure, the source IP address, and the date and time.

After you successfully log in and then log out, the failed login counter resets to zero so you will see new failed login details, if any, the next time you log in.

STEP 2 | Locate hosts that are continually attempting to login to your firewall or Panorama management server.

1. Click the failed login caution symbol to view the failed login attempts summary.
2. Locate and record the source IP address of the host that attempted to log in. For example, the following figure shows multiple failed login attempts from the IP address 192.168.2.10.



3. Work with your network administrator to locate the user and host that is using the IP address that you identified.

   If you cannot locate the system that is performing the brute-force attack, consider renaming the account to prevent future attacks.

STEP 3 | Take the following actions if you detect an account compromise.

1. Select **Monitor** > **Logs** > **Configuration** and view the configuration changes and commit history to determine if your account was used to make changes without your knowledge.
2. Select **Device** > **Config Audit** to compare the current configuration and the configuration that was running just prior to the configuration you suspect was changed using your credentials. You can also do this using Panorama.

If your administrator account was used to create a new account, performing a configuration audit helps you detect changes that are associated with any unauthorized accounts, as well.

3. Revert the configuration to a known good configuration if you see that logs were deleted or if you have difficulty determining if improper changes were made using your account.

Before you commit to a previous configuration, review it to ensure that it contains the correct settings. For example, the configuration that you revert to may not contain recent changes, so apply those changes after you commit the backup configuration.

© 2018 Palo Alto Networks, Inc.

Use the following best practices to help prevent brute-force attacks on privileged accounts.

- *Limit the number of failed attempts allowed before the firewall locks a privileged account by setting the number of Failed Attempts and the Lockout Time (min) in the authentication profile or in the Authentication Settings for the Management interface (Device > Setup > Management > Authentication Settings).*
- Use Interface Management Profiles to Restrict Access.
- *Enforce* complex passwords *for privileged accounts.*

Manage and Monitor Administrative Tasks

The Task Manager displays details about all the operations that you and other administrators initiated (such as manual commits) or that the firewall initiated (such as scheduled report generation) since the last firewall reboot. You can use the Task Manager to troubleshoot failed operations, investigate warnings associated with completed commits, view details about queued commits, or cancel pending commits.

You can also view System Logs *to monitor system events on the firewall or view* Config Logs *to monitor firewall configuration changes.*

STEP 1 | Click **Tasks** at the bottom of the web interface.

STEP 2 | **Show** only **Running** tasks (in progress) or **All** tasks (default). Optionally, filter the tasks by type:

- **Jobs**—Administrator-initiated commits, firewall-initiated commits, and software or content downloads and installations.
- **Reports**—Scheduled reports.
- **Log Requests**—Log queries that you trigger by accessing the **Dashboard** or a **Monitor** page.

STEP 3 | Perform any of the following actions:

- **Display or hide task details**—By default, the Task Manager displays the Type, Status, Start Time, and Messages for each task. To see the End Time and Job ID for a task, you must manually configure the display to expose those columns. To display or hide a column, open the drop-down in any column header, select **Columns**, and select or deselect the column names as needed.
- **Investigate warnings or failures**—Read the entries in the Messages column for task details. If the column says `Too many messages`, click the corresponding entry in the Type column to see more information.
- **Display a commit description**—If an administrator entered a description when configuring a commit, you can click **Commit Description** in the Messages column to display the description.
- **Check the position of a commit in the queue**—The Messages column indicates the queue position of commits that are in progress.
- **Cancel pending commits**—Click **Clear Commit Queue** to cancel all pending commits (available only to predefined administrative roles). To cancel an individual commit, click **x** in the Action column for that commit (the commit remains in the queue until the firewall dequeues it). You cannot cancel commits that are in progress.

Commit, Validate, and Preview Firewall Configuration Changes

A commit is the process of activating pending changes to the firewall configuration. You can filter pending changes by administrator or *location* and then preview, validate, or commit only those changes. The locations can be specific virtual systems, shared policies and objects, or shared device and network settings.

© 2018 Palo Alto Networks, Inc.

The firewall queues commit requests so that you can initiate a new commit while a previous commit is in progress. The firewall performs the commits in the order they are initiated but prioritizes auto-commits that are initiated by the firewall (such as FQDN refreshes). However, if the queue already has the maximum number of administrator-initiated commits, you must wait for the firewall to finish processing a pending commit before initiating a new one. To cancel pending commits or view details about commits of any status, see Manage and Monitor Administrative Tasks.

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



The commit, validate, preview, save, and revert operations apply only to changes made after the last commit. To restore configurations to the state they were in before the last commit, you must load a previously backed up configuration.

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

STEP 1 | Configure the scope of configuration changes that you will commit, validate, or preview.

1. Click **Commit** at the top of the web interface.
2. Select one of the following options:
   - **Commit All Changes** (default)—Applies the commit to all changes for which you have administrative privileges. You cannot manually filter the commit scope when you select this option. Instead, the administrator role assigned to the account you used to log in determines the commit scope.
   - **Commit Changes Made By**—Enables you to filter the commit scope by administrator or location. The administrative role assigned to the account you used to log in determines which changes you can filter.

   To commit the changes of other administrators, the account you used to log in must be assigned the Superuser role or an Admin Role profile *with the Commit For Other Admins privilege enabled.*
3. (Optional) To filter the commit scope by administrator, select **Commit Changes Made By**, click the adjacent link, select the administrators, and click **OK**.
4. (Optional) To filter by location, select **Commit Changes Made By** and clear any changes that you want to exclude from the Commit Scope.

   If dependencies between the configuration changes you included and excluded cause a validation error, perform the commit with all the changes included. For example, when you commit changes to a virtual system, you must include the changes of all administrators who added, deleted, or repositioned rules for the same rulebase in that virtual system.

STEP 2 | Preview the changes that the commit will activate.

This can be useful if, for example, you don't remember all your changes and you're not sure you want to activate all of them.

The firewall compares the configurations you selected in the Commit Scope to the running configuration. The preview window displays the configurations side-by-side and uses color coding to indicate which changes are additions (green), modifications (yellow), or deletions (red).

© 2018 Palo Alto Networks, Inc.

Preview Changes and select the **Lines of Context**, which is the number of lines from the compared configuration files to display before and after each highlighted difference. These additional lines help you correlate the preview output to settings in the web interface. Close the preview window when you finish reviewing the changes.

Because the preview results display in a new browser window, your browser must allow pop-ups. If the preview window does not open, refer to your browser documentation for the steps to allow pop-ups.

STEP 3 | Preview the individual settings for which you are committing changes.

This can be useful if you want to know details about the changes, such as the types of settings and who changed them.

1. Click **Change Summary**.
2. (Optional) **Group By** a column name (such as the **Type** of setting).
3. **Close** the Change Summary dialog when you finish reviewing the changes.

STEP 4 | Validate the changes before you commit to ensure the commit will succeed.

1. **Validate Changes**.

   The results display all the errors and warnings that an actual commit would display.
2. Resolve any errors that the validation results identify.

STEP 5 | Commit your configuration changes.

Commit your changes to validate and activate them.

To view details about commits that are pending (which you can still cancel), in progress, completed, or failed, see Manage and Monitor Administrative Tasks*.*

Use Global Find to Search the Firewall or Panorama Management Server

Global Find enables you to search the candidate configuration on a firewall or on Panorama for a particular string, such as an IP address, object name, policy rule name, threat ID, application name. In addition to searching for configuration objects and settings, you can search by job ID or job type for manual commits that administrators performed or auto-commits that the firewall or Panorama performed. The search results are grouped by category and provide links to the configuration location in the web interface, so that you can easily find all of the places where the string is referenced. The search results also help you identify other objects that depend on or make reference to the search term or string. For example, when deprecating a security profile enter the profile name in Global Find to locate all instances of the profile and then click each instance to navigate to the configuration page and make the necessary change. After all references are removed, you can then delete the profile. You can do this for any configuration item that has dependencies.

Watch the video.

Global Find will not search dynamic content (such as logs, address ranges, or allocated DHCP addresses). In the case of DHCP, you can search on a DHCP server attribute, such as the DNS entry, but you cannot search for individual addresses allocated to users. Global

© 2018 Palo Alto Networks, Inc.

Find also does not search for individual user or group names identified by User-ID unless the user/group is defined in a policy. In general, you can only search content that the firewall writes to the configuration.

- Launch Global Find by clicking the **Search** icon located on the upper right of the web interface.



- To access the Global Find from within a configuration area, click the drop-down next to an item and select **Global Find**:



For example, click **Global Find** on a zone named **l3-vlan-trust** to search the candidate configuration for each location where the zone is referenced. The following screen capture shows the search results for the zone l3-vlan-trust:



Search tips:

- If you initiate a search on a firewall that has multiple virtual systems enabled or if custom Administrative Role Types are defined, Global Find will only return results for areas of the firewall in which the administrator has permissions. The same applies to Panorama device groups.
- Spaces in search terms are handled as AND operations. For example, if you search on **`corp policy`**, the search results include instances where corp and policy exist in the configuration.
- To find an exact phrase, enclose the phrase in quotation marks.
- To rerun a previous search, click Search (located on the upper right of the web interface) to see a list of the last 20 searches. Click an item in the list to rerun that search. Search history is unique to each administrator account.

Manage Locks for Restricting Configuration Changes

You can use configuration locks to prevent other administrators from changing the candidate configuration or from committing configuration changes until you manually remove the lock or the firewall automatically removes it (after a commit). Locks ensure that administrators don't make conflicting changes to the same settings or interdependent settings during concurrent login sessions.

© 2018 Palo Alto Networks, Inc.

The firewall queues commit requests and performs them in the order that administrators initiate the commits. For details, see Commit, Validate, and Preview Firewall Configuration Changes. *To view the status of queued commits, see* Manage and Monitor Administrative Tasks.

- View details about current locks.

  For example, you can check whether other administrators have set locks and read comments they entered to explain the locks.

  Click the lock at the top of the web interface. An adjacent number indicates the number of current locks.

- Lock a configuration.
  1. Click the lock at the top of the web interface.

     The lock image varies based on whether existing locks are or are not set.
  2. **Take a Lock** and select the lock **Type**:
     - **Config**—Blocks other administrators from changing the candidate configuration.
     - **Commit**—Blocks other administrators from committing changes made to the candidate configuration.
  3. (Firewall with multiple virtual systems only) Select a **Location** to lock the configuration for a specific virtual system or the **Shared** location.
  4. (Optional) As a best practice, enter a **Comment** so that other administrators will understand the reason for the lock.
  5. Click **OK** and **Close**.

- Unlock a configuration.

  Only a superuser or the administrator who locked the configuration can manually unlock it. However, the firewall automatically removes a lock after completing the commit operation.

  1. Click the lock at the top of the web interface.
  2. Select the lock entry in the list.
  3. Click **Remove Lock**, **OK**, and **Close**.

- Configure the firewall to automatically apply a commit lock when you change the candidate configuration. This setting applies to all administrators.
  1. Select **Device** > **Setup** > **Management** and edit the General Settings.
  2. Select **Automatically Acquire Commit Lock** and then click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Manage Configuration Backups

The running configuration on the firewall comprises all settings you have committed and that are therefore active, such as policy rules that currently block or allow various types of traffic in your network. The candidate configuration is a copy of the running configuration plus any inactive changes that you made after the last commit. Saving backup versions of the running or candidate configuration enables you to later restore those versions. For example, if a commit validation shows that the current candidate configuration has more errors than you want to fix, you can restore a previous candidate configuration. You can also revert to the current running configuration without saving a backup first.



See Commit, Validate, and Preview Firewall Configuration Changes *for details about commit operations.*

- Save and Export Firewall Configurations
- Revert Firewall Configuration Changes

Save and Export Firewall Configurations

Saving a backup of the candidate configuration to persistent storage on the firewall enables you to later revert to that backup (see Revert Firewall Configuration Changes). This is useful for preserving changes that would otherwise be lost if a system event or administrator action causes the firewall to reboot. After rebooting, PAN-OS automatically reverts to the current version of the running configuration, which the firewall stores in a file named running-config.xml. Saving backups is also useful if you want to revert to a firewall configuration that is earlier than the current running configuration. The firewall does not automatically save the candidate configuration to persistent storage. You must manually save the candidate configuration as a default snapshot file (.snapshot.xml) or as a custom-named snapshot file. The firewall stores the snapshot file locally but you can export it to an external host.



You don't have to save a configuration backup to revert the changes made since the last commit or reboot; just select Config > Revert Changes (see Revert Firewall Configuration Changes*).*

When you edit a setting and click OK, the firewall updates the candidate configuration but does not save a backup snapshot.

Additionally, saving changes does not activate them. To activate changes, perform a commit (see Commit, Validate, and Preview Firewall Configuration Changes*).*

Palo Alto Networks recommends that you back up any important configuration to a host external to the firewall.

STEP 1 | Save a local backup snapshot of the candidate configuration if it contains changes that you want to preserve in the event the firewall reboots.

These are changes you are not ready to commit—for example, changes you cannot finish in the current login session.

To overwrite the default snapshot file (.snapshot.xml) with all the changes that all administrators made, perform one of the following steps:

- Select **Device** > **Setup** > **Operations** and **Save candidate configuration**.

© 2018 Palo Alto Networks, Inc.

- Log in to the firewall with an administrative account that is assigned the Superuser role or an Admin Role profile with the **Save For Other Admins** privilege enabled. Then select **Config** > **Save Changes** at the top of the web interface, select **Save All Changes** and **Save**.

To create a snapshot that includes all the changes that all administrators made but without overwriting the default snapshot file:

1. Select **Device** > **Setup** > **Operations** and **Save named configuration snapshot**.
2. Specify the **Name** of a new or existing configuration file.
3. Click **OK** and **Close**.

To save only specific changes to the candidate configuration without overwriting any part of the default snapshot file:

1. Log in to the firewall with an administrative account that has the role privileges required to save the desired changes.
2. Select **Config** > **Save Changes** at the top of the web interface.
3. Select **Save Changes Made By**.
4. To filter the Save Scope by administrator, click **<administrator-name>**, select the administrators, and click **OK**.
5. To filter the Save Scope by location, clear any locations that you want to exclude. The locations can be specific virtual systems, shared policies and objects, or shared device and network settings.
6. Click **Save**, specify the **Name** of a new or existing configuration file, and click **OK**.

STEP 2 | Export a candidate configuration, a running configuration, or the firewall state information to a host external to the firewall.

Select **Device** > **Setup** > **Operations** and click an export option:

- **Export named configuration snapshot**—Export the current running configuration, a named candidate configuration snapshot, or a previously imported configuration (candidate or running). The firewall exports the configuration as an XML file with the **Name** you specify.
- **Export configuration version**—Select a **Version** of the running configuration to export as an XML file. The firewall creates a version whenever you commit configuration changes.
- **Export device state**—Export the firewall state information as a bundle. Besides the running configuration, the state information includes device group and template settings pushed from Panorama. If the firewall is a GlobalProtect portal, the information also includes certificate information, a list of satellites, and satellite authentication information. If you replace a firewall or portal, you can restore the exported information on the replacement by importing the state bundle.

Revert Firewall Configuration Changes

Revert operations replace settings in the current candidate configuration with settings from another configuration. Reverting changes is useful when you want to undo changes to multiple settings as a single operation instead of manually reconfiguring each setting.

You can revert pending changes that were made to the firewall configuration since the last commit. The firewall provides the option to filter the pending changes by administrator or *location*. The locations can be specific virtual systems, shared policies and objects, or shared device and network settings. If you saved a snapshot file for a candidate configuration that is earlier than the current running configuration (see Save and Export Firewall Configurations), you can also revert to that snapshot. Reverting to a snapshot enables you to restore a candidate configuration that existed before the last commit. The firewall automatically saves a new version of the running configuration whenever you commit changes, and you can restore any of those versions.

- Revert to the current running configuration (file named running-config.xml).

© 2018 Palo Alto Networks, Inc.

This operation undoes changes you made to the candidate configuration since the last commit.

To revert all the changes that all administrators made, perform one of the following steps:

- Select **Device** > **Setup** > **Operations, Revert to running configuration**, and click **Yes** to confirm the operation.
- Log in to the firewall with an administrative account that is assigned the Superuser role or an Admin Role profile with the **Commit For Other Admins** privilege enabled. Then select **Config** > **Revert Changes** at the top of the web interface, select **Revert All Changes** and **Revert**.

To revert only specific changes to the candidate configuration:

1. Log in to the firewall with an administrative account that has the role privileges required to revert the desired changes.

   The privileges that control commit operations also control revert operations.
2. Select **Config** > **Revert Changes** at the top of the web interface.
3. Select **Revert Changes Made By**.
4. To filter the Revert Scope by administrator, click **<administrator-name>**, select the administrators, and click **OK**.
5. To filter the Revert Scope by location, clear any locations that you want to exclude.
6. **Revert** the changes.

- Revert to the default snapshot of the candidate configuration.

  This is the snapshot that you create or overwrite when you click **Config** > **Save Changes** at the top of the web interface.

  1. Select **Device** > **Setup** > **Operations** and **Revert to last saved configuration**.
  2. Click **Yes** to confirm the operation.
  3. (Optional) Click **Commit** to overwrite the running configuration with the snapshot.

- Revert to a previous version of the running configuration that is stored on the firewall.

  The firewall creates a version whenever you commit configuration changes.

  1. Select **Device** > **Setup** > **Operations** and **Load configuration version**.
  2. Select a configuration **Version** and click **OK**.
  3. (Optional) Click **Commit** to overwrite the running configuration with the version you just restored.

- Revert to one of the following:
  - Custom-named version of the running configuration that you previously imported
  - Custom-named candidate configuration snapshot (instead of the default snapshot)

  1. Select **Device** > **Setup** > **Operations** and click **Load named configuration snapshot**.
  2. Select the snapshot **Name** and click **OK**.
  3. (Optional) Click **Commit** to overwrite the running configuration with the snapshot.

- Revert to a running or candidate configuration that you previously exported to an external host.

  1. Select **Device** > **Setup** > **Operations**, click **Import named configuration snapshot**, **Browse** to the configuration file on the external host, and click **OK**.
  2. Click **Load named configuration snapshot**, select the **Name** of the configuration file you just imported, and click **OK**.

© 2018 Palo Alto Networks, Inc.

3. (Optional) Click **Commit** to overwrite the running configuration with the snapshot you just imported.

- Restore state information that you exported from a firewall.

  Besides the running configuration, the state information includes device group and template settings pushed from Panorama. If the firewall is a GlobalProtect portal, the information also includes certificate information, a list of satellites, and satellite authentication information. If you replace a firewall or portal, can you can restore the information on the replacement by importing the state bundle.

  Import state information:

  1. Select **Device** > **Setup** > **Operations**, click **Import device state**, **Browse** to the state bundle, and click **OK**.
  2. (Optional) Click **Commit** to apply the imported state information to the running configuration.

© 2018 Palo Alto Networks, Inc.

Manage Firewall Administrators

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



As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This enables you to better protect the firewall from unauthorized configuration and enables logging of the actions of individual administrators.

Make sure you are following the Best Practices for Securing Administrative Access *to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.*

- Administrative Role Types
- Configure an Admin Role Profile
- Administrative Authentication
- Configure Administrative Accounts and Authentication

Administrative Role Types

A *role* defines the type of access that an administrator has to the firewall. The role types are:

- **Admin Role Profiles**—Custom roles you can configure for more granular access control over the functional areas of the web interface, CLI, and XML API. For example, you can create an Admin Role profile for your operations staff that provides access to the firewall and network configuration areas of the web interface and a separate profile for your security administrators that provides access to security policy definitions, logs, and reports. On a firewall with multiple virtual systems, you can select whether the role defines access for all virtual systems or specific virtual systems. When new features are added to the product, you must update the roles with corresponding access privileges: the firewall does not automatically add new features to custom role definitions. For details on the privileges you can configure for custom administrator roles, see Reference: Web Interface Administrator Access.
- **Dynamic Roles**—These are built-in roles that provide access to the firewall. When new features are added, the firewall automatically updates the definitions of dynamic roles; you never need to manually update them. The following table lists the access privileges associated with dynamic roles.

Dynamic Role	Privileges
Superuser	Full access to the firewall, including defining new administrator accounts and virtual systems. You must have superuser privileges to create an administrative user with superuser privileges.
Superuser (read-only)	Read-only access to the firewall.
Virtual system administrator	Full access to a selected virtual system (vsys) on the firewall.
Virtual system administrator (read-only)	Read-only access to a selected vsys on the firewall.
Device administrator	Full access to all firewall settings except for defining new accounts or virtual systems.

© 2018 Palo Alto Networks, Inc.

Dynamic Role	Privileges
Device administrator (read-only)	Read-only access to all firewall settings except password profiles (no access) and administrator accounts (only the logged in account is visible).

Configure an Admin Role Profile

Admin Role profiles enable you to define granular administrative access privileges to ensure protection for sensitive company information and privacy for end users.



As a best practice, create Admin Role profiles that allow administrators to access only the areas of the management interfaces that they need to access to perform their jobs.

STEP 1 | Select **Device** > **Admin Roles** and click **Add**.

STEP 2 | Enter a **Name** to identify the role.

STEP 3 | For the scope of the **Role**, select **Device** or **Virtual System**.

STEP 4 | In the **Web UI** and **XML API** tabs, click the icon for each functional area to toggle it to the desired setting: Enable, Read Only, or Disable. For details on the **Web UI** options, see Web Interface Access Privileges.

STEP 5 | Select the **Command Line** tab and select a CLI access option. The **Role** scope controls the available options:

- **Device** role—**superuser**, **superreader**, **deviceadmin**, **devicereader**, or **None**
- **Virtual System** role—**vsysadmin**, **vsysreader**, or **None**

STEP 6 | Click **OK** to save the profile.

STEP 7 | Assign the role to an administrator. See Configure a Firewall Administrator Account.

Administrative Authentication

You can configure the following types of authentication and authorization (role and access domain assignment) for firewall administrators:

Authentication Method	Authorization Method	Description
Local	Local	The administrative account credentials and authentication mechanisms are local to the firewall. You can define the accounts with or without a user database that is local to the firewall—see Local Authentication for the advantages and disadvantages of using a local database. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure Local or External Authentication for Firewall Administrators.

© 2018 Palo Alto Networks, Inc.

Authentication Method	Authorization Method	Description
SSH Keys	Local	The administrative accounts are local to the firewall, but authentication to the CLI is based on SSH keys. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure SSH Key-Based Administrator Authentication to the CLI.
Certificates	Local	The administrative accounts are local to the firewall, but authentication to the web interface is based on client certificates. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure Certificate-Based Administrator Authentication to the Web Interface.
External service	Local	The administrative accounts you define locally on the firewall serve as references to the accounts defined on an external Multi-Factor Authentication, SAML, Kerberos, TACACS+, RADIUS, or LDAP server. The external server performs authentication. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure Local or External Authentication for Firewall Administrators.
External service	External service	The administrative accounts are defined on an external SAML, TACACS+, or RADIUS server. The server performs both authentication and authorization. For authorization, you define Vendor-Specific Attributes (VSAs) on the TACACS+ or RADIUS server, or SAML attributes on the SAML server. PAN-OS maps the attributes to administrator roles, access domains, user groups, and virtual systems that you define on the firewall. For details, see: • Configure SAML Authentication • Configure TACACS+ Authentication • Configure RADIUS Authentication

Configure Administrative Accounts and Authentication

If you have already configured an authentication profile (see Configure an Authentication Profile and Sequence) or you don't require one to authenticate administrators, you are ready to Configure a Firewall Administrator Account. Otherwise, perform one of the other procedures listed below to configure administrative accounts for specific types of authentication.

- Configure a Firewall Administrator Account
- Configure Local or External Authentication for Firewall Administrators
- Configure Certificate-Based Administrator Authentication to the Web Interface
- Configure SSH Key-Based Administrator Authentication to the CLI

Configure a Firewall Administrator Account

Administrative accounts specify roles and authentication methods for firewall administrators. The service that you use to assign roles and perform authentication determines whether you add the accounts on the firewall, on an external server, or both (see Administrative Authentication). If the authentication method relies on a local firewall database or an external service, you must configure an authentication profile before adding an administrative account (see Configure Administrative Accounts and Authentication). If

© 2018 Palo Alto Networks, Inc.

you already configured the authentication profile or you will use Local Authentication without a firewall database, perform the following steps to add an administrative account on the firewall.



As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This enables you to better protect the firewall from unauthorized configuration and enables logging of the actions of individual administrators.

Make sure you are following the Best Practices for Securing Administrative Access *to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.*

STEP 1 | Select **Device** > **Administrators** and **Add** an account.

STEP 2 | Enter a user **Name**.

If the firewall uses a local user database to authenticate the account, enter the name that you specified for the account in the database (see Add the user group to the local database.)

STEP 3 | Select an **Authentication Profile** or sequence if you configured either for the administrator.

If the firewall uses Local Authentication without a local user database for the account, select **None** (default) and enter a **Password**.

STEP 4 | Select the **Administrator Type**.

If you configured a custom role for the user, select **Role Based** and select the Admin Role **Profile**. Otherwise, select **Dynamic** (default) and select a dynamic role. If the dynamic role is **virtual system administrator**, add one or more virtual systems that the virtual system administrator is allowed to manage.

STEP 5 | (Optional) Select a **Password Profile** for administrators that the firewall authenticates locally without a local user database. For details, see Define a Password Profile.

STEP 6 | Click **OK** and **Commit**.

Configure Local or External Authentication for Firewall Administrators

You can use Local Authentication and External Authentication Services to authenticate administrators who access the firewall. These authentication methods prompt administrators to respond to one or more authentication challenges, such as a login page for entering a username and password.



If you use an external service to manage both authentication and authorization (role and access domain assignments), see:

- Configure SAML Authentication
- Configure TACACS+ Authentication
- Configure RADIUS Authentication

To authenticate administrators without a challenge-response mechanism, you can Configure Certificate-Based Administrator Authentication to the Web Interface *and* Configure SSH Key-Based Administrator Authentication to the CLI.

STEP 1 | (External authentication only) Enable the firewall to connect to an external server for authenticating administrators.

© 2018 Palo Alto Networks, Inc.

Configure a server profile:

- Add a RADIUS server profile.

  If the firewall integrates with a Multi-Factor Authentication (MFA) service through RADIUS, you must add a RADIUS server profile. In this case, the MFA service provides all the authentication factors (challenges). If the firewall integrates with an MFA service through a vendor API, you can still use a RADIUS server profile for the first factor but MFA server profiles are required for additional factors.
- Add an MFA server profile.
- Add a TACACS+ server profile.
- Add a SAML IdP server profile. You cannot combine Kerberos single sign-on (SSO) with SAML SSO; you can use only one type of SSO service.
- Add a Kerberos server profile.
- Add an LDAP server profile.

STEP 2 | (Local database authentication only) Configure a user database that is local to the firewall.

1. Add the user account to the local database.
2. (Optional) Add the user group to the local database.

STEP 3 | (Local authentication only) Define password complexity and expiration settings.

These settings help protect the firewall against unauthorized access by making it harder for attackers to guess passwords.

1. Define global password complexity and expiration settings for all local administrators. The settings don't apply to local database accounts for which you specified a password hash instead of a password (see Local Authentication).
   1. Select **Device** > **Setup** > **Management** and edit the Minimum Password Complexity settings.
   2. Select **Enabled**.
   3. Define the password settings and click **OK**.
2. Define a Password Profile.

   You assign the profile to administrator accounts for which you want to override the global password expiration settings. The profiles are available only to accounts that are not associated with a local database (see Local Authentication).
   1. Select **Device** > **Password Profiles** and **Add** a profile.
   2. Enter a **Name** to identify the profile.
   3. Define the password expiration settings and click **OK**.

STEP 4 | (Kerberos SSO only) Create a Kerberos keytab.

A keytab is a file that contains Kerberos account information for the firewall. To support Kerberos SSO, your network must have a Kerberos infrastructure.

STEP 5 | Configure an authentication profile.



If your administrative accounts are stored across multiple types of servers, you can create an authentication profile for each type and add all the profiles to an authentication sequence.

Configure an Authentication Profile and Sequence. In the authentication profile, specify the **Type** of authentication service and related settings:

- **External service**—Select the **Type** of external service and select the **Server Profile** you created for it.
- **Local database authentication**—Set the **Type** to **Local Database**.

© 2018 Palo Alto Networks, Inc.

- **Local authentication without a database**—Set the **Type** to **None**.
- **Kerberos SSO**—Specify the **Kerberos Realm** and **Import** the **Kerberos Keytab**.

STEP 6 | Assign the authentication profile or sequence to an administrator account.

1. Configure a Firewall Administrator Account.
   - Assign the **Authentication Profile** or sequence that you configured.
   - (Local database authentication only) Specify the **Name** of the user account you added to the local database.
2. **Commit** your changes.
3. (Optional) Test Authentication Server Connectivity to verify that the firewall can use the authentication profile to authenticate administrators.

Configure Certificate-Based Administrator Authentication to the Web Interface

As a more secure alternative to password-based authentication to the firewall web interface, you can configure certificate-based authentication for administrator accounts that are local to the firewall. Certificate-based authentication involves the exchange and verification of a digital signature instead of a password.

Configuring certificate-based authentication for any administrator disables the username/password logins for all administrators on the firewall; administrators thereafter require the certificate to log in.

STEP 1 | Generate a certificate authority (CA) certificate on the firewall.

You will use this CA certificate to sign the client certificate of each administrator.

Create a Self-Signed Root CA Certificate.

Alternatively, Import a Certificate and Private Key *from your enterprise CA or a third-party CA.*

STEP 2 | Configure a certificate profile for securing access to the web interface.

Configure a Certificate Profile.

- Set the **Username Field** to **Subject**.
- In the CA Certificates section, **Add** the **CA Certificate** you just created or imported.

STEP 3 | Configure the firewall to use the certificate profile for authenticating administrators.

1. Select **Device** > **Setup** > **Management** and edit the Authentication Settings.
2. Select the **Certificate Profile** you created for authenticating administrators and click **OK**.

STEP 4 | Configure the administrator accounts to use client certificate authentication.

For each administrator who will access the firewall web interface, Configure a Firewall Administrator Account and select **Use only client certificate authentication**.

If you have already deployed client certificates that your enterprise CA generated, skip to Step 8. Otherwise, go to Step 5.

STEP 5 | Generate a client certificate for each administrator.

Generate a Certificate. In the **Signed By** drop-down, select a self-signed root CA certificate.

© 2018 Palo Alto Networks, Inc.

STEP 6 | Export the client certificate.

1. Export a Certificate and Private Key.
2. **Commit** your changes. The firewall restarts and terminates your login session. Thereafter, administrators can access the web interface only from client systems that have the client certificate you generated.

STEP 7 | Import the client certificate into the client system of each administrator who will access the web interface.

Refer to your web browser documentation.

STEP 8 | Verify that administrators can access the web interface.

1. Open the firewall IP address in a browser on the computer that has the client certificate.
2. When prompted, select the certificate you imported and click **OK**. The browser displays a certificate warning.
3. Add the certificate to the browser exception list.
4. Click **Login**. The web interface should appear without prompting you for a username or password.

Configure SSH Key-Based Administrator Authentication to the CLI

For administrators who use Secure Shell (SSH) to access the CLI of a Palo Alto Networks firewall, SSH keys provide a more secure authentication method than passwords. SSH keys almost eliminate the risk of brute-force attacks, provide the option for two-factor authentication (key and passphrase), and don't send passwords over the network. SSH keys also enable automated scripts to access the CLI.

STEP 1 | Use an SSH key generation tool to create an asymmetric keypair on the client system of the administrator.

The supported key formats are IETF SECSH and Open SSH. The supported algorithms are DSA (1,024 bits) and RSA (768-4,096 bits).

For the commands to generate the keypair, refer to your SSH client documentation.

The public key and private key are separate files. Save both to a location that the firewall can access. For added security, enter a passphrase to encrypt the private key. The firewall prompts the administrator for this passphrase during login.

STEP 2 | Configure the administrator account to use public key authentication.

1. Configure a Firewall Administrator Account.
   - Configure the authentication method to use as a fallback if SSH key authentication fails. If you configured an **Authentication Profile** for the administrator, select it in the drop-down. If you select **None**, you must enter a **Password** and **Confirm Password**.
   - Select **Use Public Key Authentication (SSH)**, then **Import Key**, **Browse** to the public key you just generated, and click **OK**.
2. **Commit** your changes.

STEP 3 | Configure the SSH client to use the private key to authenticate to the firewall.

Perform this task on the client system of the administrator. For the steps, refer to your SSH client documentation.

STEP 4 | Verify that the administrator can access the firewall CLI using SSH key authentication.

1. Use a browser on the client system of the administrator to go to the firewall IP address.
2. Log in to the firewall CLI as the administrator. After entering a username, you will see the following output (the key value is an example):

© 2018 Palo Alto Networks, Inc.

```
Authenticating with public key “dsa-key-20130415”
```

3. If prompted, enter the passphrase you defined when creating the keys.

© 2018 Palo Alto Networks, Inc.

Reference: Web Interface Administrator Access

You can configure privileges for an entire firewall or for one or more virtual systems (on platforms that support multiple virtual systems). Within that **Device** or **Virtual System** designation, you can configure privileges for custom administrator roles, which are more granular than the fixed privileges associated with a dynamic administrator role.

Configuring privileges at a granular level ensures that lower level administrators cannot access certain information. You can create custom roles for firewall administrators (see Configure a Firewall Administrator Account), Panorama administrators, or Device Group and Template administrators (refer to the Panorama Administrator's Guide). You apply the admin role to a custom role-based administrator account where you can assign one or more virtual systems. The following topics describe the privileges you can configure for custom administrator roles.

- Web Interface Access Privileges
- Panorama Web Interface Access Privileges

Web Interface Access Privileges

If you want to prevent a role-based administrator from accessing specific tabs on the web interface, you can disable the tab and the administrator will not even see it when logging in using the associated role-based administrative account. For example, you could create an Admin Role Profile for your operations staff that provides access to the **Device** and **Network** tabs only and a separate profile for your security administrators that provides access to the **Object**, **Policy**, and **Monitor** tabs.

An admin role can apply at the **Device** level or **Virtual System** level as defined by the **Device** or **Virtual System** radio button. If you select **Virtual System**, the admin assigned this profile is restricted to the virtual system(s) he or she is assigned to. Furthermore, only the **Device** > **Setup** > **Services** > **Virtual Systems** tab is available to that admin, not the **Global** tab.

The following topics describe how to set admin role privileges to the different parts of the web interface:

- Define Access to the Web Interface Tabs
- Provide Granular Access to the Monitor Tab
- Provide Granular Access to the Policy Tab
- Provide Granular Access to the Objects Tab
- Provide Granular Access to the Network Tab
- Provide Granular Access to the Device Tab
- Define User Privacy Settings in the Admin Role Profile
- Restrict Administrator Access to Commit and Validate Functions
- Provide Granular Access to Global Settings
- Provide Granular Access to the Panorama Tab

Define Access to the Web Interface Tabs

The following table describes the top-level access privileges you can assign to an admin role profile (**Device** > **Admin Roles**). You can enable, disable, or define read-only access privileges at the top-level tabs in the web interface.

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Dashboard	Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets.	Yes	No	Yes
ACC	Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy** > **Show Full IP Addresses** option and/or the **Show User Names In Logs And Reports** option.	Yes	No	Yes
Monitor	Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the administrator can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab.	Yes	No	Yes
Policies	Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the administrator can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab.	Yes	No	Yes
Objects	Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the administrator can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab.	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Network	Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the administrator can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab.	Yes	No	Yes
Device	Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any firewall-wide configuration information, such as User-ID, high availability, server profile or certificate configuration information. For more granular control over what objects the administrator can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab. *You cannot enable access to the Admin Roles or Administrators nodes for a role-based administrator even if you enable full access to the Device tab.*	Yes	No	Yes

Provide Granular Access to the Monitor Tab

In some cases you might want to enable the administrator to view some but not all areas of the **Monitor** tab. For example, you might want to restrict operations administrators to the Config and System logs only, because they do not contain sensitive user data. Although this section of the administrator role definition specifies what areas of the **Monitor** tab the administrator can see, you can also couple privileges in this section with privacy privileges, such as disabling the ability to see usernames in logs and reports. One thing to keep in mind, however, is that any system-generated reports will still show usernames and IP addresses even if you disable that functionality in the role. For this reason, if you do not want the administrator to see any of the private user information, disable access to the specific reports as detailed in the following table.

The following table lists the **Monitor** tab access levels and the administrator roles for which they are available.

Device Group and Template roles can see log data only for the device groups that are within the access domains assigned to those roles.

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Monitor	Enables or disables access to the **Monitor** tab. If disabled, the administrator will not see this tab or any of the associated logs or reports.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Logs	Enables or disables access to all log files. You can also leave this privilege enabled and then disable specific logs that you do not want the administrator to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the logs, you can disable the **Privacy** > **Show Full IP Addresses** option and/or the **Show User Names In Logs And Reports** option.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Traffic	Specifies whether the administrator can see the traffic logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Threat	Specifies whether the administrator can see the threat logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
URL Filtering	Specifies whether the administrator can see the URL filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
WildFire Submissions	Specifies whether the administrator can see the WildFire logs. These logs are only available if you have a WildFire subscription.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Data Filtering	Specifies whether the administrator can see the data filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
HIP Match	Specifies whether the administrator can see the HIP Match logs. HIP	Firewall: Yes	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	Match logs are only available if you have a GlobalProtect portal license and gateway subscription.	Panorama: Yes Device Group/Template: Yes			
User-ID	Specifies whether the administrator can see the User-ID logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Tunnel Inspection	Specifies whether the administrator can see the Tunnel Inspection logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Configuration	Specifies whether the administrator can see the configuration logs.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
System	Specifies whether the administrator can see the system logs.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Alarms	Specifies whether the administrator can see system-generated alarms.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Authentication	Specifies whether the administrator can see the Authentication logs.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Automated Correlation Engine	Enables or disables access to the correlation objects and correlated event logs generated on the firewall.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Correlation Objects	Specifies whether the administrator can view and enable/disable the correlation objects.	Firewall: Yes Panorama: Yes	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
		Device Group/Template: Yes			
Correlated Events	Specifies whether the administrator	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Packet Capture	Specifies whether the administrator can see packet captures (pcaps) from the **Monitor** tab. Keep in mind that packet captures are raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the pcap and you should therefore disable the Packet Capture privilege if you are concerned about user privacy.	Firewall: Yes Panorama: No Device Group/Template: No	Yes	Yes	Yes
App Scope	Specifies whether the administrator can see the App Scope visibility and analysis tools. Enabling App Scope enables access to all of the **App Scope** charts.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Session Browser	Specifies whether the administrator can browse and filter current running sessions on the firewall. Keep in mind that the session browser shows raw flow data and as such may contain user IP addresses. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in the session browser and you should therefore disable the **Session Browser** privilege if you are concerned about user privacy.	Firewall: Yes Panorama: No Device Group/Template: No	Yes	No	Yes
Block IP List	Specifies whether the administrator can view the block list (Enable or Read Only) and delete entries from the list (Enable). If you disable the setting, the administrator won't be able to view or delete entries from the block list.	Firewall: Yes Panorama: under Context Switch UI: Yes Template: Yes	Yes	Yes	Yes
Botnet	Specifies whether the administrator can generate and view botnet	Firewall: Yes	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	analysis reports or view botnet reports in read-only mode. Disabling the **Show Full IP Addresses** privileges will not obfuscate the IP address in scheduled botnet reports and you should therefore disable the **Botnet** privilege if you are concerned about user privacy.	Panorama: No Device Group/Template: No			
PDF Reports	Enables or disables access to all PDF reports. You can also leave this privilege enabled and then disable specific PDF reports that you do not want the administrator to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy** > **Show Full IP Addresses** option and/or the **Show User Names In Logs And Reports** option.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Manage PDF Summary	Specifies whether the administrator can view, add or delete PDF summary report definitions. With read-only access, the administrator can see PDF summary report definitions, but not add or delete them. If you disable this option, the administrator can neither view the report definitions nor add/delete them.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
PDF Summary Reports	Specifies whether the administrator can see the generated PDF Summary reports in **Monitor** > **Reports**. If you disable this option, the **PDF Summary Reports** category will not display in the **Reports** node.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
User Activity Report	Specifies whether the administrator can view, add or delete User Activity report definitions and download the reports. With read-only access, the administrator can see User Activity report definitions, but not add, delete, or download them. If you disable this option,	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	the administrator cannot see this category of PDF report.				
SaaS Application Usage Report	Specifies whether the administrator can view, add or delete a SaaS application usage report. With read-only access, the administrator can see the SaaS application usage report definitions, but cannot add or delete them. If you disable this option, the administrator can neither view the report definitions nor add or delete them.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
Report Groups	Specifies whether the administrator can view, add or delete report group definitions. With read-only access, the administrator can see report group definitions, but not add or delete them. If you disable this option, the administrator cannot see this category of PDF report.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
Email Scheduler	Specifies whether the administrator can schedule report groups for email. Because the generated reports that get emailed may contain sensitive user data that is not removed by disabling the **Privacy** > **Show Full IP Addresses** option and/or the **Show User Names In Logs And Reports** options and because they may also show log data to which the administrator does not have access, you should disable the **Email Scheduler** option if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
Manage Custom Reports	Enables or disables access to all custom report functionality. You can also leave this privilege enabled and then disable specific custom report categories that you do not want the administrator to be able to access. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the **Privacy** > **Show Full IP Addresses** option and/or the **Show**	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	User Names In Logs And Reports option. *Reports that are scheduled to run rather than run on demand will show IP address and user information. In this case, be sure to restrict access to the corresponding report areas. In addition, the custom report feature does not restrict the ability to generate reports that contain log data contained in logs that are excluded from the administrator role.*				
Application Statistics	Specifies whether the administrator can create a custom report that includes data from the application statistics database.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Data Filtering Log	Specifies whether the administrator can create a custom report that includes data from the Data Filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Threat Log	Specifies whether the administrator can create a custom report that includes data from the Threat logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Threat Summary	Specifies whether the administrator can create a custom report that includes data from the Threat Summary database.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Traffic Log	Specifies whether the administrator can create a custom report that includes data from the Traffic logs.	Firewall: Yes Panorama: Yes	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
		Device Group/Template: Yes			
Traffic Summary	Specifies whether the administrator can create a custom report that includes data from the Traffic Summary database.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
URL Log	Specifies whether the administrator can create a custom report that includes data from the URL Filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Hipmatch	Specifies whether the administrator can create a custom report that includes data from the HIP Match logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
WildFire Log	Specifies whether the administrator can create a custom report that includes data from the WildFire logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Userid	Specifies whether the administrator can create a custom report that includes data from the User-ID logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Auth	Specifies whether the administrator can create a custom report that includes data from the Authentication logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
View Scheduled Custom Reports	Specifies whether the administrator can view a custom report that has been scheduled to generate.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
View Predefined Application Reports	Specifies whether the administrator can view Application Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	reports if you have user privacy requirements.				
View Predefined Threat Reports	Specifies whether the administrator can view Threat Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
View Predefined URL Filtering Reports	Specifies whether the administrator can view URL Filtering Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
View Predefined Traffic Reports	Specifies whether the administrator can view Traffic Reports. Privacy privileges do not impact reports available on the **Monitor** > **Reports** node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes

Provide Granular Access to the Policy Tab

If you enable the Policy option in the Admin Role profile, you can then enable, disable, or provide read-only access to specific nodes within the tab as necessary for the role you are defining. By enabling access to a specific policy type, you enable the ability to view, add, or delete policy rules. By enabling read-only access to a specific policy, you enable the administrator to view the corresponding policy rule base, but not add or delete rules. Disabling access to a specific type of policy prevents the administrator from seeing the policy rule base.

Because policy that is based on specific users (by user name or IP address) must be explicitly defined, privacy settings that disable the ability to see full IP addresses or user names do not apply to the Policy tab. Therefore, you should only allow access to the Policy tab to administrators that are excluded from user privacy restrictions.

Access Level	Description	Enable	Read Only	Disable
Security	Enable this privilege to allow the administrator to view, add, and/or delete security rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	from seeing the security rulebase, disable this privilege.			
NAT	Enable this privilege to allow the administrator to view, add, and/or delete NAT rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the NAT rulebase, disable this privilege.	Yes	Yes	Yes
QoS	Enable this privilege to allow the administrator to view, add, and/or delete QoS rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the QoS rulebase, disable this privilege.	Yes	Yes	Yes
Policy Based Forwarding	Enable this privilege to allow the administrator to view, add, and/or delete Policy-Based Forwarding (PBF) rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the PBF rulebase, disable this privilege.	Yes	Yes	Yes
Decryption	Enable this privilege to allow the administrator to view, add, and/or delete decryption rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the decryption rulebase, disable this privilege.	Yes	Yes	Yes
Tunnel Inspection	Enable this privilege to allow the administrator to view, add, and/or delete Tunnel Inspection rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the Tunnel Inspection rulebase, disable this privilege.	Yes	Yes	Yes
Application Override	Enable this privilege to allow the administrator to view, add, and/or delete application override policy rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the application override rulebase, disable this privilege.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Authentication	Enable this privilege to allow the administrator to view, add, and/or delete Authentication policy rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the Authentication rulebase, disable this privilege.	Yes	Yes	Yes
DoS Protection	Enable this privilege to allow the administrator to view, add, and/or delete DoS protection rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the DoS protection rulebase, disable this privilege.	Yes	Yes	Yes

Provide Granular Access to the Objects Tab

An *object* is a container that groups specific policy filter values—such as IP addresses, URLs, applications, or services—for simplified rule definition. For example, an address object might contain specific IP address definitions for the web and application servers in your DMZ zone.

When deciding whether to allow access to the objects tab as a whole, determine whether the administrator will have policy definition responsibilities. If not, the administrator probably does not need access to the tab. If, however, the administrator will need to create policy, you can enable access to the tab and then provide granular access privileges at the node level.

By enabling access to a specific node, you give the administrator the privilege to view, add, and delete the corresponding object type. Giving read-only access allows the administrator to view the already defined objects, but not create or delete any. Disabling a node prevents the administrator from seeing the node in the web interface.

Access Level	Description	Enable	Read Only	Disable
Addresses	Specifies whether the administrator can view, add, or delete address objects for use in security policy.	Yes	Yes	Yes
Address Groups	Specifies whether the administrator can view, add, or delete address group objects for use in security policy.	Yes	Yes	Yes
Regions	Specifies whether the administrator can view, add, or delete regions objects for use in security, decryption, or DoS policy.	Yes	Yes	Yes
Applications	Specifies whether the administrator can view, add, or delete application objects for use in policy.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Application Groups	Specifies whether the administrator can view, add, or delete application group objects for use in policy.	Yes	Yes	Yes
Application Filters	Specifies whether the administrator can view, add, or delete application filters for simplification of repeated searches.	Yes	Yes	Yes
Services	Specifies whether the administrator can view, add, or delete service objects for use in creating policy rules that limit the port numbers an application can use.	Yes	Yes	Yes
Service Groups	Specifies whether the administrator can view, add, or delete service group objects for use in security policy.	Yes	Yes	Yes
Tags	Specifies whether the administrator can view, add, or delete tags that have been defined on the firewall.	Yes	Yes	Yes
GlobalProtect	Specifies whether the administrator can view, add, or delete HIP objects and profiles. You can restrict access to both types of objects at the GlobalProtect level, or provide more granular control by enabling the GlobalProtect privilege and restricting HIP Object or HIP Profile access.	Yes	No	Yes
HIP Objects	Specifies whether the administrator can view, add, or delete HIP objects, which are used to define HIP profiles. HIP Objects also generate HIP Match logs.	Yes	Yes	Yes
Clientless Apps	Specifies whether the administrator can view, add, modify, or delete GlobalProtect VPN Clientless applications.	Yes	Yes	Yes
Clientless App Groups	Specifies whether the administrator can view, add, modify, or delete GlobalProtect VPN Clientless application groups.	Yes	Yes	Yes
HIP Profiles	Specifies whether the administrator can view, add, or delete HIP Profiles for use in security policy and/or for generating HIP Match logs.	Yes	Yes	Yes
External Dynamic Lists	Specifies whether the administrator can view, add, or delete dynamic block lists for use in security policy.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Custom Objects	Specifies whether the administrator can see the custom spyware and vulnerability signatures. You can restrict access to either enable or disable access to all custom signatures at this level, or provide more granular control by enabling the Custom Objects privilege and then restricting access to each type of signature.	Yes	No	Yes
Data Patterns	Specifies whether the administrator can view, add, or delete custom data pattern signatures for use in creating custom Vulnerability Protection profiles.	Yes	Yes	Yes
Spyware	Specifies whether the administrator can view, add, or delete custom spyware signatures for use in creating custom Vulnerability Protection profiles.	Yes	Yes	Yes
Vulnerability	Specifies whether the administrator can view, add, or delete custom vulnerability signatures for use in creating custom Vulnerability Protection profiles.	Yes	Yes	Yes
URL Category	Specifies whether the administrator can view, add, or delete custom URL categories for use in policy.	Yes	Yes	Yes
Security Profiles	Specifies whether the administrator can see security profiles. You can restrict access to either enable or disable access to all security profiles at this level, or provide more granular control by enabling the Security Profiles privilege and then restricting access to each type of profile.	Yes	No	Yes
Antivirus	Specifies whether the administrator can view, add, or delete antivirus profiles.	Yes	Yes	Yes
Anti-Spyware	Specifies whether the administrator can view, add, or delete Anti-Spyware profiles.	Yes	Yes	Yes
Vulnerability Protection	Specifies whether the administrator can view, add, or delete Vulnerability Protection profiles.	Yes	Yes	Yes
URL Filtering	Specifies whether the administrator can view, add, or delete URL filtering profiles.	Yes	Yes	Yes
File Blocking	Specifies whether the administrator can view, add, or delete file blocking profiles.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
WildFire Analysis	Specifies whether the administrator can view, add, or delete WildFire analysis profiles.	Yes	Yes	Yes
Data Filtering	Specifies whether the administrator can view, add, or delete data filtering profiles.	Yes	Yes	Yes
DoS Protection	Specifies whether the administrator can view, add, or delete DoS protection profiles.	Yes	Yes	Yes
Security Profile Groups	Specifies whether the administrator can view, add, or delete security profile groups.	Yes	Yes	Yes
Log Forwarding	Specifies whether the administrator can view, add, or delete log forwarding profiles.	Yes	Yes	Yes
Authentication	Specifies whether the administrator can view, add, or delete authentication enforcement objects.	Yes	Yes	Yes
Decryption Profile	Specifies whether the administrator can view, add, or delete decryption profiles.	Yes	Yes	Yes
Schedules	Specifies whether the administrator can view, add, or delete schedules for limiting a security policy to a specific date and/or time range.	Yes	Yes	Yes

Provide Granular Access to the Network Tab

When deciding whether to allow access to the **Network** tab as a whole, determine whether the administrator will have network administration responsibilities, including GlobalProtect administration. If not, the administrator probably does not need access to the tab.

You can also define access to the **Network** tab at the node level. By enabling access to a specific node, you give the administrator the privilege to view, add, and delete the corresponding network configurations. Giving read-only access allows the administrator to view the already-defined configuration, but not create or delete any. Disabling a node prevents the administrator from seeing the node in the web interface.

Access Level	Description	Enable	Read Only	Disable
Interfaces	Specifies whether the administrator can view, add, or delete interface configurations.	Yes	Yes	Yes
Zones	Specifies whether the administrator can view, add, or delete zones.	Yes	Yes	Yes
VLANs	Specifies whether the administrator can view, add, or delete VLANs.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Virtual Wires	Specifies whether the administrator can view, add, or delete virtual wires.	Yes	Yes	Yes
Virtual Routers	Specifies whether the administrator can view, add, modify or delete virtual routers.	Yes	Yes	Yes
IPSec Tunnels	Specifies whether the administrator can view, add, modify, or delete IPSec Tunnel configurations.	Yes	Yes	Yes
DHCP	Specifies whether the administrator can view, add, modify, or delete DHCP server and DHCP relay configurations.	Yes	Yes	Yes
DNS Proxy	Specifies whether the administrator can view, add, modify, or delete DNS proxy configurations.	Yes	Yes	Yes
GlobalProtect	Specifies whether the administrator can view, add, modify GlobalProtect portal and gateway configurations. You can disable access to the GlobalProtect functions entirely, or you can enable the GlobalProtect privilege and then restrict the role to either the portal or gateway configuration areas.	Yes	No	Yes
Portals	Specifies whether the administrator can view, add, modify, or delete GlobalProtect portal configurations.	Yes	Yes	Yes
Gateways	Specifies whether the administrator can view, add, modify, or delete GlobalProtect gateway configurations.	Yes	Yes	Yes
MDM	Specifies whether the administrator can view, add, modify, or delete GlobalProtect MDM server configurations.	Yes	Yes	Yes
Device Block List	Specifies whether the administrator can view, add, modify, or delete device block lists.	Yes	Yes	Yes
Clientless Apps	Specifies whether the administrator can view, add, modify, or delete GlobalProtect Clientless VPN applications.	Yes	Yes	Yes
Clientless App Groups	Specifies whether the administrator can view, add, modify, or delete GlobalProtect Clientless VPN application groups.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
QoS	Specifies whether the administrator can view, add, modify, or delete QoS configurations.	Yes	Yes	Yes
LLDP	Specifies whether the administrator can view add, modify, or delete LLDP configurations.	Yes	Yes	Yes
Network Profiles	Sets the default state to enable or disable for all of the Network settings described below.	Yes	No	Yes
GlobalProtect IPSec Crypto	Controls access to the **Network Profiles** > **GlobalProtect IPSec Crypto** node. If you disable this privilege, the administrator will not see that node, or configure algorithms for authentication and encryption in VPN tunnels between a GlobalProtect gateway and clients. If you set the privilege to read-only, the administrator can view existing GlobalProtect IPSec Crypto profiles but cannot add or edit them.	Yes	Yes	Yes
IKE Gateways	Controls access to the **Network Profiles** > **IKE Gateways** node. If you disable this privilege, the administrator will not see the **IKE Gateways** node or define gateways that include the configuration information necessary to perform IKE protocol negotiation with peer gateway. If the privilege state is set to read-only, you can view the currently configured IKE Gateways but cannot add or edit gateways.	Yes	Yes	Yes
IPSec Crypto	Controls access to the **Network Profiles** > **IPSec Crypto** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **IPSec Crypto** node or specify protocols and algorithms for identification, authentication, and encryption in VPN tunnels based on IPSec SA negotiation. If the privilege state is set to read-only, you can view the currently configured IPSec Crypto configuration but cannot add or edit a configuration.	Yes	Yes	Yes
IKE Crypto	Controls how devices exchange information to ensure secure communication. Specify the protocols and algorithms for identification, authentication, and encryption in VPN	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	tunnels based on IPsec SA negotiation (IKEv1 Phase-1).			
Monitor	Controls access to the **Network Profiles** > **Monitor** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **Monitor** node or be able to create or edit a monitor profile that is used to monitor IPSec tunnels and monitor a next-hop device for policy-based forwarding (PBF) rules. If the privilege state is set to read-only, you can view the currently configured monitor profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes
Interface Mgmt	Controls access to the **Network Profiles** > **Interface Mgmt** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **Interface Mgmt** node or be able to specify the protocols that are used to manage the firewall. If the privilege state is set to read-only, you can view the currently configured Interface management profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes
Zone Protection	Controls access to the **Network Profiles** > **Zone Protection** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **Zone Protection** node or be able to configure a profile that determines how the firewall responds to attacks from specified security zones. If the privilege state is set to read-only, you can view the currently configured Zone Protection profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes
QoS Profile	Controls access to the **Network Profiles** > **QoS** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **QoS** node or be able to configure a QoS profile that determines how QoS traffic classes are treated. If the privilege state is set to read-only, you can view the currently configured QoS profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
LLDP Profile	Controls access to the **Network Profiles** > **LLDP** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **LLDP** node or be able to configure an LLDP profile that controls whether the interfaces on the firewall can participate in the Link Layer Discovery Protocol. If the privilege state is set to read-only, you can view the currently configured LLDP profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes
BFD Profile	Controls access to the **Network Profiles** > **BFD Profile** node. If you disable this privilege, the administrator will not see the **Network Profiles** > **BFD Profile** node or be able to configure a BFD profile. A Bidirectional Forwarding Detection (BFD) profile allows you to configure BFD settings to apply to one or more static routes or routing protocols. Thus, BFD detects a failed link or BFD peer and allows an extremely fast failover. If the privilege state is set to read-only, you can view the currently configured BFD profile but cannot add or edit a BFD profile.	Yes	Yes	Yes

Provide Granular Access to the Device Tab

To define granular access privileges for the **Device** tab, when creating or editing an admin role profile (**Device** > **Admin Roles**), scroll down to the **Device** node on the **WebUI** tab.

Access Level	Description	Enable	Read Only	Disable
Setup	Controls access to the **Setup** node. If you disable this privilege, the administrator will not see the **Setup** node or have access to firewall-wide setup configuration information, such as Management, Operations, Service, Content-ID, Wildfire or Session setup information. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
Management	Controls access to the **Management** node. If you disable this privilege, the administrator will not be able to configure settings such as the hostname, domain, timezone,	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	authentication, logging and reporting, Panorama connections, banner, message, and password complexity settings, and more. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.			
Operations	Controls access to the **Operations** and **Telemetry and Threat Intelligence** nodes. If you disable this privilege, the administrator cannot: • Load firewall configurations. • Save or revert the firewall configuration.  *This privilege applies only to the Device > Operations options. The Save and Commit privileges control whether the administrator can save or revert configurations through the Config > Save and Config > Revert options.* • Create custom logos. • Configure SNMP monitoring of firewall settings. • Configure the Statistics Service feature. • Configure **Telemetry and Threat Intelligence** settings. Only administrators with the predefined Superuser role can export or import firewall configurations and shut down the firewall. Only administrators with the predefined Superuser or Device Administrator role can reboot the firewall or restart the dataplane. Administrators with a role that allows access only to specific virtual systems cannot load, save, or revert firewall configurations through the **Device** > **Operations** options.	Yes	Yes	Yes
Services	Controls access to the **Services** node. If you disable this privilege, the administrator will not be able to configure services for DNS servers, an update server, proxy server, or NTP servers, or set up service routes.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.			
Content-ID	Controls access to the **Content-ID** node. If you disable this privilege, the administrator will not be able to configure URL filtering or Content-ID. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
WildFire	Controls access to the **WildFire** node. If you disable this privilege, the administrator will not be able to configure WildFire settings. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
Session	Controls access to the **Session** node. If you disable this privilege, the administrator will not be able to configure session settings or timeouts for TCP, UDP or ICMP, or configure decryption or VPN session settings. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
HSM	Controls access to the **HSM** node. If you disable this privilege, the administrator will not be able to configure a Hardware Security Module. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
High Availability	Controls access to the **High Availability** node. If you disable this privilege, the administrator will not see the **High Availability** node or have access to firewall-wide high availability configuration information such as General setup information or Link and Path Monitoring. If you set this privilege to read-only, the administrator can view High Availability configuration information for the firewall but is not allowed to perform any configuration procedures.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Config Audit	Controls access to the **Config Audit** node. If you disable this privilege, the administrator will not see the **Config Audit** node or have access to any firewall-wide configuration information.	Yes	No	Yes
Administrators	Controls access to the **Administrators** node. This function can only be allowed for read-only access. If you disable this privilege, the administrator will not see the **Administrators** node or have access to information about their own administrator account. If you set this privilege to read-only, the administrator can view the configuration information for their own administrator account. They will not see any information about other administrator accounts configured on the firewall.	No	Yes	Yes
Admin Roles	Controls access to the **Admin Roles** node. This function can only be allowed for read-only access. If you disable this privilege, the administrator will not see the **Admin Roles** node or have access to any firewall-wide information concerning Admin Role profiles configuration. If you set this privilege to read-only, you can view the configuration information for all administrator roles configured on the firewall.	No	Yes	Yes
Authentication Profile	Controls access to the **Authentication Profile** node. If you disable this privilege, the administrator will not see the **Authentication Profile** node or be able to create or edit authentication profiles that specify RADIUS, TACACS+, LDAP, Kerberos, SAML, multi-factor authentication (MFA), or local database authentication settings. PAN-OS uses authentication profiles to authenticate firewall administrators and Captive Portal or GlobalProtect end users. If you set this privilege to read-only, the administrator can view the **Authentication Profile** information but cannot create or edit authentication profiles.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Authentication Sequence	Controls access to the **Authentication Sequence** node. If you disable this privilege, the administrator will not see the **Authentication Sequence** node or be able to create or edit an authentication sequence. If you set this privilege to read-only, the administrator can view the **Authentication Profile** information but cannot create or edit an authentication sequence.	Yes	Yes	Yes
Virtual Systems	Controls access to the **Virtual Systems** node. If you disable this privilege, the administrator will not see or be able to configure virtual systems. If the privilege state is set to read-only, you can view the currently configured virtual systems but cannot add or edit a configuration.	Yes	Yes	Yes
Shared Gateways	Controls access to the **Shared Gateways** node. Shared gateways allow virtual systems to share a common interface for external communications. If you disable this privilege, the administrator will not see or be able to configure shared gateways. If the privilege state is set to read-only, you can view the currently configured shared gateways but cannot add or edit a configuration.	Yes	Yes	Yes
User Identification	Controls access to the **User Identification** node. If you disable this privilege, the administrator will not see the **User Identification** node or have access to firewall-wide User Identification configuration information, such as User Mapping, Connection Security, User-ID Agents, Terminal Services Agents, Group Mappings Settings, or Captive Portal Settings. If you set this privilege to read-only, the administrator can view configuration information for the firewall but is not allowed to perform any configuration procedures.	Yes	Yes	Yes
VM Information Source	Controls access to the **VM Information Source** node that allows you to configure the firewall/Windows User-ID agent to collect VM inventory automatically. If you disable this	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	privilege, the administrator will not see the **VM Information Source** node. If you set this privilege to read-only, the administrator can view the VM information sources configured but cannot add, edit, or delete any sources. *This privilege is not available to Device Group and Template administrators.*			
Certificate Management	Sets the default state to enable or disable for all of the Certificate settings described below.	Yes	No	Yes
Certificates	Controls access to the **Certificates** node. If you disable this privilege, the administrator will not see the **Certificates** node or be able to configure or access information regarding Device Certificates or Default Trusted Certificate Authorities. If you set this privilege to read-only, the administrator can view Certificate configuration information for the firewall but is not allowed to perform any configuration procedures.	Yes	Yes	Yes
Certificate Profile	Controls access to the **Certificate Profile** node. If you disable this privilege, the administrator will not see the **Certificate Profile** node or be able to create certificate profiles. If you set this privilege to read-only, the administrator can view Certificate Profiles that are currently configured for the firewall but is not allowed to create or edit a certificate profile.	Yes	Yes	Yes
OCSP Responder	Controls access to the **OCSP Responder** node. If you disable this privilege, the administrator will not see the **OCSP Responder** node or be able to define a server that will be used to verify the revocation status of certificates issues by the firewall. If you set this privilege to read-only, the administrator can view the **OCSP Responder** configuration for the firewall but is not allowed to create or edit an OCSP responder configuration.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
SSL/TLS Service Profile	Controls access to the **SSL/TLS Service Profile** node. If you disable this privilege, the administrator will not see the node or configure a profile that specifies a certificate and a protocol version or range of versions for firewall services that use SSL/TLS. If you set this privilege to read-only, the administrator can view existing SSL/TLS Service profiles but cannot create or edit them.	Yes	Yes	Yes
SCEP	Controls access to the **SCEP** node. If you disable this privilege, the administrator will not see the node or be able to define a profile that specifies simple certificate enrollment protocol (SCEP) settings for issuing unique device certificates. If you set this privilege to read-only, the administrator can view existing SCEP profiles but cannot create or edit them.	Yes	Yes	Yes
SSL Decryption Exclusion	Controls access to the **SSL Decryption Exclusion** node. If you disable this privilege, the administrator will not see the node or be able see the SSL decryption add custom exclusions. If you set this privilege to read-only, the administrator can view existing SSL decryption exceptions but cannot create or edit them.	Yes	Yes	Yes
Response Pages	Controls access to the **Response Pages** node. If you disable this privilege, the administrator will not see the **Response Page** node or be able to define a custom HTML message that is downloaded and displayed instead of a requested web page or file. If you set this privilege to read-only, the administrator can view the **Response Page** configuration for the firewall but is not allowed to create or edit a response page configuration.	Yes	Yes	Yes
Log Settings	Sets the default state to enable or disable for all of the Log settings described below.	Yes	No	Yes
System	Controls access to the **Log Settings** > **System** node. If you disable this privilege, the administrator cannot see the **Log Settings** >	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	System node or specify which System logs the firewall forwards to Panorama or external services (such as a syslog server). If you set this privilege to read-only, the administrator can view the **Log Settings** > **System** settings for the firewall but cannot add, edit, or delete the settings.			
Configuration	Controls access to the **Log Settings** > **Configuration** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **Configuration** node or specify which Configuration logs the firewall forwards to Panorama or external services (such as a syslog server). If you set this privilege to read-only, the administrator can view the **Log Settings** > **Configuration** settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
User-ID	Controls access to the **Log Settings** > **User-ID** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **User-ID** node or specify which User-ID logs the firewall forwards to Panorama or external services (such as a syslog server). If you set this privilege to read-only, the administrator can view the **Log Settings** > **User-ID** settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
HIP Match	Controls access to the **Log Settings** > **HIP Match** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **HIP Match** node or specify which Host Information Profile (HIP) match logs the firewall forwards to Panorama or external services (such as a syslog server). HIP match logs provide information on Security policy rules that apply to GlobalProtect clients If you set this privilege to read-only, the administrator can view the **Log Settings** > **HIP** settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
Correlation	Controls access to the **Log Settings** > **Correlation** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **Correlation** node or add, delete, or	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	modify correlation log forwarding settings or tag source or destination IP addresses. If you set this privilege to read-only, the administrator can view the **Log Settings** > **Correlation** settings for the firewall but cannot add, edit, or delete the settings.			
Alarm Settings	Controls access to the **Log Settings** > **Alarm Settings** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **Alarm Settings** node or configure notifications that the firewall generates when a Security policy rule (or group of rules) is hit repeatedly within a configurable time period. If you set this privilege to read-only, the administrator can view the **Log Settings** > **Alarm Settings** for the firewall but cannot edit the settings.	Yes	Yes	Yes
Manage Logs	Controls access to the **Log Settings** > **Manage Logs** node. If you disable this privilege, the administrator cannot see the **Log Settings** > **Manage Logs** node or clear the indicated logs. If you set this privilege to read-only, the administrator can view the **Log Settings** > **Manage Logs** information but cannot clear any of the logs.	Yes	Yes	Yes
Server Profiles	Sets the default state to enable or disable for all of the Server Profiles settings described below.	Yes	No	Yes
SNMP Trap	Controls access to the **Server Profiles** > **SNMP Trap** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **SNMP Trap** node or be able to specify one or more SNMP trap destinations to be used for system log entries. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **SNMP Trap Logs** information but cannot specify SNMP trap destinations.	Yes	Yes	Yes
Syslog	Controls access to the **Server Profiles** > **Syslog** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **Syslog** node or be able to specify one or more syslog servers.	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	If you set this privilege to read-only, the administrator can view the **Server Profiles** > **Syslog** information but cannot specify syslog servers.			
Email	Controls access to the **Server Profiles** > **Email** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **Email** node or be able to configure an email profile that can be used to enable email notification for system and configuration log entries. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **Email** information but cannot configure an email server profile.	Yes	Yes	Yes
HTTP	Controls access to the **Server Profiles** > **HTTP** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **HTTP** node or be able to configure an HTTP server profile that can be used to enable log forwarding to HTTP destinations any log entries. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **HTTP** information but cannot configure an HTTP server profile.	Yes	Yes	Yes
Netflow	Controls access to the **Server Profiles** > **Netflow** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **Netflow** node or be able to define a NetFlow server profile, which specifies the frequency of the export along with the NetFlow servers that will receive the exported data. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **Netflow** information but cannot define a Netflow profile.	Yes	Yes	Yes
RADIUS	Controls access to the **Server Profiles** > **RADIUS** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **RADIUS** node or be able to configure settings for the RADIUS servers that are identified in authentication profiles. If you set this privilege to read-only, the administrator can view the **Server Profiles** >	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	RADIUS information but cannot configure settings for the RADIUS servers.			
TACACS+	Controls access to the **Server Profiles** > **TACACS+** node. If you disable this privilege, the administrator will not see the node or configure settings for the TACACS+ servers that authentication profiles reference. If you set this privilege to read-only, the administrator can view existing TACACS+ server profiles but cannot add or edit them.	Yes	Yes	Yes
LDAP	Controls access to the **Server Profiles** > **LDAP** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **LDAP** node or be able to configure settings for the LDAP servers to use for authentication by way of authentication profiles. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **LDAP** information but cannot configure settings for the LDAP servers.	Yes	Yes	Yes
Kerberos	Controls access to the **Server Profiles** > **Kerberos** node. If you disable this privilege, the administrator will not see the **Server Profiles** > **Kerberos** node or configure a Kerberos server that allows users to authenticate natively to a domain controller. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **Kerberos** information but cannot configure settings for Kerberos servers.	Yes	Yes	Yes
SAML Identity Provider	Controls access to the **Server Profiles** > **SAML Identity Provider** node. If you disable this privilege, the administrator cannot see the node or configure SAML identity provider (IdP) server profiles. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **SAML Identity Provider** information but cannot configure SAML IdP server profiles.	Yes	Yes	Yes
Multi Factor Authentication	Controls access to the **Server Profiles** > **Multi Factor Authentication** node. If you disable this privilege, the administrator cannot see the			

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	node or configure multi-factor authentication (MFA) server profiles. If you set this privilege to read-only, the administrator can view the **Server Profiles** > **SAML Identity Provider** information but cannot configure MFA server profiles.			
Local User Database	Sets the default state to enable or disable for all of the Local User Database settings described below.	Yes	No	Yes
Users	Controls access to the **Local User Database** > **Users** node. If you disable this privilege, the administrator will not see the **Local User Database** > **Users** node or set up a local database on the firewall to store authentication information for remote access users, firewall administrators, and Captive Portal users. If you set this privilege to read-only, the administrator can view the **Local User Database** > **Users** information but cannot set up a local database on the firewall to store authentication information.	Yes	Yes	Yes
User Groups	Controls access to the **Local User Database** > **Users** node. If you disable this privilege, the administrator will not see the **Local User Database** > **Users** node or be able to add user group information to the local database. If you set this privilege to read-only, the administrator can view the **Local User Database** > **Users** information but cannot add user group information to the local database.	Yes	Yes	Yes
Access Domain	Controls access to the **Access Domain** node. If you disable this privilege, the administrator will not see the **Access Domain** node or be able to create or edit an access domain. If you set this privilege to read-only, the administrator can view the **Access Domain** information but cannot create or edit an access domain.	Yes	Yes	Yes
Scheduled Log Export	Controls access to the **Scheduled Log Export** node. If you disable this privilege, the administrator will not see the **Scheduled Log Export** node or be able schedule exports of logs and save them to a File Transfer Protocol (FTP) server in CSV format or use Secure Copy	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	(SCP) to securely transfer data between the firewall and a remote host. If you set this privilege to read-only, the administrator can view the **Scheduled Log Export Profile** information but cannot schedule the export of logs.			
Software	Controls access to the **Software** node. If you disable this privilege, the administrator will not see the **Software** node or view the latest versions of the PAN-OS software available from Palo Alto Networks, read the release notes for each version, and select a release to download and install. If you set this privilege to read-only, the administrator can view the **Software** information but cannot download or install software.	Yes	Yes	Yes
GlobalProtect Client	Controls access to the **GlobalProtect Client** node. If you disable this privilege, the administrator will not see the **GlobalProtect Client** node or view available GlobalProtect releases, download the code or activate the GlobalProtect agent. If you set this privilege to read-only, the administrator can view the available **GlobalProtect Client** releases but cannot download or install the agent software.	Yes	Yes	Yes
Dynamic Updates	Controls access to the **Dynamic Updates** node. If you disable this privilege, the administrator will not see the **Dynamic Updates** node or be able to view the latest updates, read the release notes for each update, or select an update to upload and install. If you set this privilege to read-only, the administrator can view the available **Dynamic Updates** releases, read the release notes but cannot upload or install the software.	Yes	Yes	Yes
Licenses	Controls access to the **Licenses** node. If you disable this privilege, the administrator will not see the **Licenses** node or be able to view the licenses installed or activate licenses. If you set this privilege to read-only, the administrator can view the installed **Licenses**,	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	but cannot perform license management functions.			
Support	Controls access to the **Support** node. If you disable this privilege, the administrator cannot see the **Support** node, activate support, or access production and security alerts from Palo Alto Networks. If you set this privilege to read-only, the administrator can see the **Support** node and access production and security alerts but cannot activate support. Only administrators with the predefined Superuser role can use the **Support** node to generate tech support files or generate and download stats dump and core files.	Yes	Yes	Yes
Master Key and Diagnostics	Controls access to the **Master Key and Diagnostics** node. If you disable this privilege, the administrator will not see the **Master Key and Diagnostics** node or be able to specify a master key to encrypt private keys on the firewall. If you set this privilege to read-only, the administrator can view the **Master Key and Diagnostics** node and view information about master keys that have been specified but cannot add or edit a new master key configuration.	Yes	Yes	Yes

Define User Privacy Settings in the Admin Role Profile

To define what private end user data an administrator has access to, when creating or editing an admin role profile (**Device** > **Admin Roles**), scroll down to the **Privacy** option on the **WebUI** tab.

Access Level	Description	Enable	Read Only	Disable
Privacy	Sets the default state to enable or disable for all of the privacy settings described below.	Yes	N/A	Yes
Show Full IP addresses	When disabled, full IP addresses obtained by traffic running through the Palo Alto firewall are not shown in logs or reports. In place of the IP addresses that are normally displayed, the relevant subnet is displayed.	Yes	N/A	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	Scheduled reports that are displayed in the interface through Monitor > Reports and reports that are sent via scheduled emails will still display full IP addresses. Because of this exception, we recommend that the following settings within the Monitor tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler.			
Show User Names in Logs and Reports	When disabled, user names obtained by traffic running through the Palo Alto Networks firewall are not shown in logs or reports. Columns where the user names would normally be displayed are empty. *Scheduled reports that are displayed in the interface through Monitor > Reports or reports that are sent via the email scheduler will still display user names. Because of this exception, we recommend that the following settings within the Monitor tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler.*	Yes	N/A	Yes
View PCAP Files	When disabled, packet capture files that are normally available within the Traffic, Threat and Data Filtering logs are not displayed.	Yes	N/A	Yes

Restrict Administrator Access to Commit and Validate Functions

To restrict access to commit (and revert), save, and validate functions when creating or editing an Admin Role profile (**Device** > **Admin Roles**), scroll down to the **Commit**, **Save**, and **Validate** options on the **WebUI** tab.

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Commit	Sets the default state to enabled or disabled for all of the commit and revert privileges described below.	Yes	N/A	Yes
Device	When disabled, an administrator cannot commit or revert changes that any administrator made to the firewall configuration, including his or her own changes.	Yes	N/A	Yes
Commit For Other Admins	When disabled, an administrator cannot commit or revert changes that other administrators made to the firewall configuration.	Yes	N/A	Yes
Save	Sets the default state to enabled or disabled for all of the save operation privileges described below.	Yes	N/A	Yes
Partial save	When disabled, an administrator cannot save changes that any administrator made to the firewall configuration, including his or her own changes.	Yes	N/A	Yes
Save For Other Admins	When disabled, an administrator cannot save changes that other administrators made to the firewall configuration.	Yes	N/A	Yes
Validate	When disabled, an administrator cannot validate a configuration.	Yes	N/A	Yes

Provide Granular Access to Global Settings

To define what global settings and administrator has access to, when creating or editing an admin role profile (**Device** > **Admin Roles**), scroll down to the **Global** option on the **WebUI** tab.

Access Level	Description	Enable	Read Only	Disable
Global	Sets the default state to enable or disable for all of the global settings described below. In effect, this setting is only for System Alarms at this time.	Yes	N/A	Yes
System Alarms	When disabled, an administrator cannot view or acknowledge alarms that are generated.	Yes	N/A	Yes

© 2018 Palo Alto Networks, Inc.

Provide Granular Access to the Panorama Tab

The following table lists the **Panorama** tab access levels and the custom Panorama administrator roles for which they are available. Firewall administrators cannot access any of these privileges.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Setup	Specifies whether the administrator can view or edit Panorama setup information, including **Management**, **Operations and Telemetry**, **Services**, Content-ID, **WildFire**, Session, or **HSM**. If you set the privilege to: • read-only, the administrator can see the information but cannot edit it. • disable this privilege, the administrator cannot see or edit the information.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
High Availability	Specifies whether the administrator can view and manage high availability (HA) settings for the Panorama management server. If you set this privilege to read-only, the administrator can view HA configuration information for the Panorama management server but can't manage the configuration. If you disable this privilege, the administrator can't see or manage HA configuration settings for the Panorama management server.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Config Audit	Specifies whether the administrator can run Panorama configuration audits. If you disable this privilege, the administrator can't run Panorama configuration audits.	Panorama: Yes Device Group/Template: No	Yes	No	Yes
Administrators	Specifies whether the administrator can view Panorama administrator account details. You can't enable full access to this function: just read-	Panorama: Yes Device Group/Template: No	No	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	only access. (Only Panorama administrators with a dynamic role can add, edit, or delete Panorama administrators.) With read-only access, the administrator can see information about his or her own account but no other Panorama administrator accounts. If you disable this privilege, the administrator can't see information about any Panorama administrator account, including his or her own.				
Admin Roles	Specifies whether the administrator can view Panorama administrator roles. You can't enable full access to this function: just read-only access. (Only Panorama administrators with a dynamic role can add, edit, or delete custom Panorama roles.) With read-only access, the administrator can see Panorama administrator role configurations but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama administrator roles.	Panorama: Yes Device Group/Template: No	No	Yes	Yes
Access Domain	Specifies whether the administrator can view, add, edit, delete, or clone access domain configurations for Panorama administrators. (This privilege controls access only to the configuration of access domains, not access to the device groups, templates, and firewall contexts that are assigned to access domains.) If you set this privilege to read-only, the administrator can view Panorama access domain configurations but can't manage them. If you disable this privilege, the administrator can't see or	Panorama: Yes Device Group/Template: No  *You assign access domains to Device Group and Template administrators so they can access the configuration and monitoring data*	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	manage Panorama access domain configurations.	*within the device groups, templates, and firewall contexts that are assigned to those access domains.*			
Authentication Profile	Specifies whether the administrator can view, add, edit, delete, or clone authentication profiles for Panorama administrators. If you set this privilege to read-only, the administrator can view Panorama authentication profiles but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama authentication profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Authentication Sequence	Specifies whether the administrator can view, add, edit, delete, or clone authentication sequences for Panorama administrators. If you set this privilege to read-only, the administrator can view Panorama authentication sequences but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama authentication sequences.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
User Identification	Specifies whether the administrator can configure User-ID connection security and view, add, edit, or delete User-ID redistribution points (such as User-ID agents). If you set this privilege to read-only, the administrator can view settings for User-ID connection	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	security and redistribution points but can't manage the settings. If you disable this privilege, the administrator can't see or manage settings for User-ID connection security or redistribution points.				
Managed Devices	Specifies whether the administrator can view, add, edit, or delete firewalls as managed devices, and install software or content updates on them. If you set this privilege to read-only, the administrator can see managed firewalls but can't add, delete, tag, or install updates on them. If you disable this privilege, the administrator can't view, add, edit, tag, delete, or install updates on managed firewalls.  *An administrator with Device Deployment privileges can still select Panorama > Device Deployment to install updates on managed firewalls.*	Panorama: Yes Device Group/Template: Yes	Yes (No for Device Group and Template roles)	Yes	Yes
Templates	Specifies whether the administrator can view, edit, add, or delete templates and template stacks. If you set the privilege to read-only, the administrator can see template and stack configurations but can't manage them. If you disable this privilege, the administrator can't see or manage template and stack configurations.	Panorama: Yes Device Group/Template: Yes  *Device Group and Template administrators can see only the templates and stacks that are within the access domains*	Yes (No for Device Group and Template admins)	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
		assigned to those administrators.			
Device Groups	Specifies whether the administrator can view, edit, add, or delete device groups. If you set this privilege to read-only, the administrator can see device group configurations but can't manage them. If you disable this privilege, the administrator can't see or manage device group configurations.	Panorama: Yes Device Group/Template: Yes  *Device Group and Template administrators can access only the device groups that are within the access domains assigned to those administrators.*	Yes	Yes	Yes
Managed Collectors	Specifies whether the administrator can view, edit, add, or delete managed collectors. If you set this privilege to read-only, the administrator can see managed collector configurations but can't manage them. If you disable this privilege, the administrator can't view, edit, add, or delete managed collector configurations.  *An administrator with Device Deployment privileges can still use the Panorama > Device Deployment options to install updates*	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	on managed collectors.				
Collector Groups	Specifies whether the administrator can view, edit, add, or delete Collector Groups. If you set this privilege to read-only, the administrator can see Collector Groups but can't manage them. If you disable this privilege, the administrator can't see or manage Collector Groups.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
VMware Service Manager	Specifies whether the administrator can view and edit VMware Service Manager settings. If you set this privilege to read-only, the administrator can see the settings but can't perform any related configuration or operational procedures. If you disable this privilege, the administrator can't see the settings or perform any related configuration or operational procedures.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Certificate Management	Sets the default state, enabled or disabled, for all of the Panorama certificate management privileges.	Panorama: Yes Device Group/Template: No	Yes	No	Yes
Certificates	Specifies whether the administrator can view, edit, generate, delete, revoke, renew, or export certificates. This privilege also specifies whether the administrator can import or export HA keys. If you set this privilege to read-only, the administrator can see Panorama certificates but can't manage the certificates or HA keys. If you disable this privilege, the administrator can't see or manage Panorama certificates or HA keys.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Certificate Profile	Specifies whether the administrator can view, add, edit, delete or clone Panorama certificate profiles. If you set this privilege to read-only, the administrator can see Panorama certificate profiles but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama certificate profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
SSL/TLS Service Profile	Specifies whether the administrator can view, add, edit, delete or clone SSL/TLS Service profiles. If you set this privilege to read-only, the administrator can see SSL/TLS Service profiles but can't manage them. If you disable this privilege, the administrator can't see or manage SSL/TLS Service profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Log Settings	Sets the default state, enabled or disabled, for all the log setting privileges.	Panorama: Yes Device Group/Template: No	Yes	No	Yes
System	Specifies whether the administrator can see and configure the settings that control the forwarding of System logs to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the System log forwarding settings but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  *This privilege pertains only to System logs that Panorama and Log Collectors generate.*	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for System logs that Log Collectors receive from firewalls. The Device > Log Settings > System privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).				
Config	Specifies whether the administrator can see and configure the settings that control the forwarding of Config logs to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the Config log forwarding settings but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  *This privilege pertains only to Config logs that Panorama and Log Collectors generate. The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for Config logs that Log Collectors receive from firewalls. The Device > Log Settings >*	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	Configuration *privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
User-ID	Specifies whether the administrator can see and configure the settings that control the forwarding of User-ID logs to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the Config log forwarding settings but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  *This privilege pertains only to User-ID logs that Panorama generates. The* Collector Groups *privilege (Panorama > Collector Groups) controls forwarding for User-ID logs that Log Collectors receive from firewalls. The Device > Log Settings >* User-ID *privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).*	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
HIP Match	Specifies whether the administrator can see and	Panorama: Yes	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	configure the settings that control the forwarding of HIP Match logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the forwarding settings of HIP Match logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  *The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for HIP Match logs that Log Collectors receive from firewalls. The Device > Log Settings > HIP Match privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).*	Device Group/Template: No			
Correlation	Specifies whether the administrator can see and configure the settings that control the forwarding of Correlation logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the Correlation log forwarding settings but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	 *The Collector Groups privilege (Panorama > Collector Groups) controls forwarding of Correlation logs from a Panorama M-Series appliance or Panorama virtual appliance in Panorama mode.*				
Traffic	Specifies whether the administrator can see and configure the settings that control the forwarding of Traffic logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the forwarding settings of Traffic logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  *The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for Traffic logs that Log Collectors receive from firewalls. The Log Forwarding privilege (Objects > Log Forwarding) controls forwarding from firewalls directly to external services (without aggregation on Log Collectors).*	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Threat	Specifies whether the administrator can see and configure the settings that control the forwarding of Threat logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the forwarding settings of Threat logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  *The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for Threat logs that Log Collectors receive from firewalls. The Log Forwarding privilege (Objects > Log Forwarding) controls forwarding from firewalls directly to external services (without aggregation on Log Collectors).*	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Wildfire	Specifies whether the administrator can see and configure the settings that control the forwarding of WildFire logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the forwarding settings of WildFire logs but can't manage them.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you disable this privilege, the administrator can't see or manage the settings. *The Collector Groups privilege (Panorama > Collector Groups) controls the forwarding for WildFire logs that Log Collectors receive from firewalls. The Log Forwarding privilege (Objects > Log Forwarding) controls forwarding from firewalls directly to external services (without aggregation on Log Collectors).*				
Server Profiles	Sets the default state, enabled or disabled, for all the server profile privileges. *These privileges pertain only to the server profiles that are used for forwarding logs from Panorama or Log Collectors and the server profiles that are used for authenticating Panorama administrators. The Device > Server Profiles privileges control access to the server profiles that are used for forwarding logs directly from firewalls to external*	Panorama: Yes Device Group/Template: No	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	services and for authenticating firewall administrators.				
SNMP Trap	Specifies whether the administrator can see and configure SNMP trap server profiles. If you set this privilege to read-only, the administrator can see SNMP trap server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage SNMP trap server profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Syslog	Specifies whether the administrator can see and configure Syslog server profiles. If you set this privilege to read-only, the administrator can see Syslog server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage Syslog server profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Email	Specifies whether the administrator can see and configure email server profiles. If you set this privilege to read-only, the administrator can see email server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage email server profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
RADIUS	Specifies whether the administrator can see and configure the RADIUS server profiles that are used to authenticate Panorama administrators. If you set this privilege to read-only, the administrator can see the	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	RADIUS server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the RADIUS server profiles.				
TACACS+	Specifies whether the administrator can see and configure the TACACS+ server profiles that are used to authenticate Panorama administrators. If you disable this privilege, the administrator can't see the node or configure settings for the TACACS + servers that authentication profiles reference. If you set this privilege to read-only, the administrator can view existing TACACS+ server profiles but can't add or edit them.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
LDAP	Specifies whether the administrator can see and configure the LDAP server profiles that are used to authenticate Panorama administrators. If you set this privilege to read-only, the administrator can see the LDAP server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the LDAP server profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Kerberos	Specifies whether the administrator can see and configure the Kerberos server profiles that are used to authenticate Panorama administrators. If you set this privilege to read-only, the administrator can see the Kerberos server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the Kerberos server profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
SAML Identity Provider	Specifies whether the administrator can see and configure the SAML Identity Provider (IdP) server profiles that are used to authenticate Panorama administrators. If you set this privilege to read-only, the administrator can see the SAML IdP server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the SAML IdP server profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Scheduled Config Export	Specifies whether the administrator can view, add, edit, delete, or clone scheduled Panorama configuration exports. If you set this privilege to read-only, the administrator can view the scheduled exports but can't manage them. If you disable this privilege, the administrator can't see or manage the scheduled exports.	Panorama: Yes Device Group/Template: No	Yes	No	Yes
Software	Specifies whether the administrator can: view information about software updates installed on the Panorama management server; download, upload, or install the updates; and view the associated release notes. If you set this privilege to read-only, the administrator can view information about Panorama software updates and view the associated release notes but can't perform any related operations. If you disable this privilege, the administrator can't see Panorama software updates, see the associated release notes, or perform any related operations.  *The Panorama* > Device Deployment >	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	Software *privilege controls access to PAN-OS software deployed on firewalls and Panorama software deployed on Dedicated Log Collectors.*				
Dynamic Updates	Specifies whether the administrator can: view information about content updates installed on the Panorama management server (for example, WildFire updates); download, upload, install, or revert the updates; and view the associated release notes. If you set this privilege to read-only, the administrator can view information about Panorama content updates and view the associated release notes but can't perform any related operations. If you disable this privilege, the administrator can't see Panorama content updates, see the associated release notes, or perform any related operations.  *The Panorama > Device Deployment > Dynamic Updates privilege controls access to content updates deployed on firewalls and Dedicated Log Collectors.*	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Support	Specifies whether the administrator can: view Panorama support license information, product alerts, and security alerts; activate a support license, and manage cases. Only a superuser	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	admin can generate Tech Support file. If you set this privilege to read-only, the administrator can view Panorama support information, product alerts, and security alerts, but can't activate a support license, generate Tech Support files, or manage cases. If you disable this privilege, the administrator can't: see Panorama support information, product alerts, or security alerts; activate a support license, generate Tech Support files, or manage cases.				
Device Deployment	Sets the default state, enabled or disabled, for all the privileges associated with deploying licenses and software or content updates to firewalls and Log Collectors.  *The Panorama > Software and Panorama > Dynamic Updates privileges control the software and content updates installed on a Panorama management server.*	Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Software	Specifies whether the administrator can: view information about the software updates installed on firewalls and Log Collectors; download, upload, or install the updates; and view the associated release notes. If you set this privilege to read-only, the administrator can see information about the software updates and view the associated release notes but can't deploy the updates to firewalls or dedicated Log Collectors.	Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you disable this privilege, the administrator can't see information about the software updates, see the associated release notes, or deploy the updates to firewalls or Dedicated Log Collectors.				
GlobalProtect Client	Specifies whether the administrator can: view information about GlobalProtect agent/app software updates on firewalls; download, upload, or activate the updates; and view the associated release notes. If you set this privilege to read-only, the administrator can see information about GlobalProtect agent/app software updates and view the associated release notes but can't activate the updates on firewalls. If you disable this privilege, the administrator can't see information about GlobalProtect agent/app software updates, see the associated release notes, or activate the updates on firewalls.	Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
Dynamic Updates	Specifies whether the administrator can: view information about the content updates (for example, Applications updates) installed on firewalls and Dedicated Log Collectors; download, upload, or install the updates; and view the associated release notes. If you set this privilege to read-only, the administrator can see information about the content updates and view the associated release notes but can't deploy the updates to firewalls or Dedicated Log Collectors. If you disable this privilege, the administrator can't see information about the content updates, see the associated release notes, or deploy the	Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	updates to firewalls or Dedicated Log Collectors.				
Licenses	Specifies whether the administrator can view, refresh, and activate firewall licenses. If you set this privilege to read-only, the administrator can view firewall licenses but can't refresh or activate those licenses. If you disable this privilege, the administrator can't view, refresh, or activate firewall licenses.	Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
Master Key and Diagnostics	Specifies whether the administrator can view and configure a master key by which to encrypt private keys on Panorama. If you set this privilege to read-only, the administrator can view the Panorama master key configuration but can't change it. If you disable this privilege, the administrator can't see or edit the Panorama master key configuration.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Panorama Web Interface Access Privileges

The custom Panorama administrator roles allow you to define access to the options on Panorama and the ability to only allow access to Device Groups and Templates (**Policies**, **Objects**, **Network**, **Device** tabs).

The administrator roles you can create are **Panorama** and **Device Group and Template**. You can't assign CLI access privileges to a **Device Group and Template** Admin Role profile. If you assign superuser privileges for the CLI to a **Panorama** Admin Role profile, administrators with that role can access all features regardless of the web interface privileges you assign.

Access Level	Description	Enable	Read Only	Disable
Dashboard	Controls access to the **Dashboard** tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets.	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
ACC	Controls access to the Application Command Center (ACC). If you disable this privilege, the **ACC** tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the **Privacy** > **Show Full IP Addresses** option and/or the **Show User Names In Logs And Reports** option.	Yes	No	Yes
Monitor	Controls access to the **Monitor** tab. If you disable this privilege, the administrator will not see the **Monitor** tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the administrator can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab.	Yes	No	Yes
Policies	Controls access to the **Policies** tab. If you disable this privilege, the administrator will not see the **Policies** tab and will not have access to any policy information. For more granular control over what policy information the administrator can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the **Policies** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab.	Yes	No	Yes
Objects	Controls access to the **Objects** tab. If you disable this privilege, the administrator will not see the **Objects** tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the administrator can see, leave the **Objects** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab.	Yes	No	Yes
Network	Controls access to the **Network** tab. If you disable this privilege, the administrator will not see the **Network** tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
	information or to the network profiles. For more granular control over what objects the administrator can see, leave the **Network** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab.			
Device	Controls access to the **Device** tab. If you disable this privilege, the administrator will not see the **Device** tab and will not have access to any firewall-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the administrator can see, leave the **Device** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab.  *You can't enable access to the Admin Roles or Administrators nodes for a role-based administrator even if you enable full access to the Device tab.*	Yes	No	Yes
Panorama	Controls access to the **Panorama** tab. If you disable this privilege, the administrator will not see the **Panorama** tab and will not have access to any Panorama-wide configuration information, such as Managed Devices, Managed Collectors, or Collector Groups. For more granular control over what objects the administrator can see, leave the **Panorama** option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Panorama Tab.	Yes	No	Yes
Privacy	Controls access to the privacy settings described in Define User Privacy Settings in the Admin Role Profile.	Yes	No	Yes
Validate	When disabled, an administrator cannot validate a configuration.	Yes	No	Yes
Save	Sets the default state (enabled or disabled) for all the save privileges described below (Partial Save and Save For Other Admins).	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
• Partial Save	When disabled, an administrator cannot save changes that any administrator made to the Panorama configuration.	Yes	No	Yes
• Save For Other Admins	When disabled, an administrator cannot save changes that other administrators made to the Panorama configuration.	Yes	No	Yes
Commit	Sets the default state (enabled or disabled) for all the commit, push, and revert privileges described below (Panorama, Device Groups, Templates, Force Template Values, Collector Groups, WildFire Appliance Clusters).	Yes	No	Yes
• Panorama	When disabled, an administrator cannot commit or revert configuration changes that any administrators made, including his or her own changes.	Yes	No	Yes
• Commit for Other Admins	When disabled, an administrator cannot commit or revert configuration changes that other administrators made.	Yes	No	Yes
Device Groups	When disabled, an administrator cannot push changes to device groups.	Yes	No	Yes
Templates	When disabled, an administrator cannot push changes to templates.	Yes	No	Yes
Force Template Values	This privilege controls access to the **Force Template Values** option in the Push Scope Selection dialog. When disabled, an administrator cannot replace overridden settings in local firewall configurations with settings that Panorama pushes from a template.	Yes	No	Yes
Collector Groups	When disabled, an administrator cannot push changes to Collector Groups.	Yes	No	Yes
WildFire Appliance Clusters	When disabled, an administrator cannot push changes to WildFire appliance clusters.	Yes	No	Yes
Tasks	When disabled, an administrator cannot access the Task Manager.	Yes	No	Yes
Global	Controls access to the global settings (system alarms) described in Provide Granular Access to Global Settings.	Yes	No	Yes

© 2018 Palo Alto Networks, Inc.

Reference: Port Number Usage

The following tables list the ports that firewalls and Panorama use to communicate with each other, or with other services on the network.

- Ports Used for Management Functions
- Ports Used for HA
- Ports Used for Panorama
- Ports Used for GlobalProtect
- Ports Used for User-ID

Ports Used for Management Functions

The firewall and Panorama use the following ports for management functions.

Destination Port	Protocol	Description
22	TCP	Used for communication from a client system to the firewall CLI interface.
80	TCP	The port the firewall listens on for Online Certificate Status Protocol (OCSP) updates when acting as an OCSP responder.
123	UDP	Port the firewall uses for NTP updates.
443	TCP	Used for communication from a client system to the firewall web interface. This is also the port the firewall and User-ID agent listens on for updates when you Enable VM Monitoring to Track Changes on the Virtual Network. For monitoring an AWS environment, this is the only port that is used. For monitoring a VMware vCenter/ESXi environment, the listening port defaults to 443, but it is configurable.
162	UDP	Port the firewall, Panorama, or a Log Collector uses to Forward Traps to an SNMP Manager. *This port doesn't need to be open on the Palo Alto Networks firewall. You must configure the Simple Network Management Protocol (SNMP) manager to listen on this port. For details, refer to the documentation of your SNMP management software.*
161	UDP	Port the firewall listens on for polling requests (GET messages) from the SNMP manager.
514 514 6514	TCP UDP SSL	Port that the firewall, Panorama, or a Log Collector uses to send logs to a syslog server if you Configure Syslog Monitoring, and the ports that the PAN-OS integrated User-ID agent or Windows-based User-ID agent listens on for authentication syslog messages.

© 2018 Palo Alto Networks, Inc.

Destination Port	Protocol	Description
2055	UDP	Default port the firewall uses to send NetFlow records to a NetFlow collector if you Configure NetFlow Exports, but this is configurable.
5008	TCP	Port the GlobalProtect Mobile Security Manager listens on for HIP requests from the GlobalProtect gateways. If you are using a third-party MDM system, you can configure the gateway to use a different port as required by the MDM vendor.
6080 6081 6082	TCP TCP TCP	Ports used for User-ID™ Captive Portal: 6080 for NT LAN Manager (NTLM) authentication, 6081 for Captive Portal in transparent mode, and 6082 for Captive Portal in redirect mode.
10443	SSL	Port that the firewall and Panorama use to provide contextual information about a threat or to seamlessly shift your threat investigation from the firewallor Panorama to the Threat Vault and AutoFocus.

Ports Used for HA

Firewalls configured as High Availability (HA) peers must be able to communicate with each other to maintain state information (HA1 control link) and synchronize data (HA2 data link). In Active/Active HA deployments the peer firewalls must also forward packets to the HA peer that owns the session. The HA3 link is a Layer 2 (MAC-in-MAC) link and it does not support Layer 3 addressing or encryption.

Destination Port	Protocol	Description
28769 28260	TCP TCP	Used for the HA1 control link for clear text communication between the HA peer firewalls. The HA1 link is a Layer 3 link and requires an IP address.
28	TCP	Used for the HA1 control link for encrypted communication (SSH over TCP) between the HA peer firewalls.
28770	TCP	Listening port for HA1 backup links.
28771	TCP	Used for heartbeat backups. Palo Alto Networks recommends enabling heartbeat backup on the MGT interface if you use an in-band port for the HA1 or the HA1 backup links.
99 29281	IP UDP	Used for the HA2 link 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 firewall (Active/Passive) or active-primary (Active/Active) to the passive firewall (Active/Passive) or active-secondary (Active/Active). The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.

© 2018 Palo Alto Networks, Inc.

Destination Port	Protocol	Description
		The HA data link can also 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.

Ports Used for Panorama

Panorama uses the following ports.

Destination Port	Protocol	Description
22	TCP	Used for communication from a client system to the Panorama CLI interface.
443	TCP	Used for communication from a client system to the Panorama web interface. Used for communication from the WF-500 Wildfire appliance to the: • Firewall • Panorama • Wildfire cloud • Content Delivery Network (CDN)
444	TCP	Used for communication between Panorama and the Logging Service and the GlobalProtect cloud service.
3978	TCP	Used for communication between Panorama and managed firewalls or managed collectors or the Logging Service, as well as for communication among managed collectors in a Collector Group: • For communication between Panorama and firewalls, this is a bi-directional connection on which the firewalls forward logs to Panorama and Panorama pushes configuration changes to the firewalls. Context switching commands are sent over the same connection. • Log Collectors use this destination port to forward logs to Panorama. • For communication with the default Log Collector on an M-Series appliance in Panorama mode and with Dedicated Log Collectors.
28443	TCP	Used for managed devices (firewalls, Wildfire appliances and Log Collectors) to retrieve software and content updates from Panorama. *Only devices that run PAN-OS 8.x and later releases retrieve updates from Panorama over this port. For devices running earlier releases,*

© 2018 Palo Alto Networks, Inc.

Destination Port	Protocol	Description
		Panorama pushes the update packages over port 3978.
10443	TCP	Used on the WildFire appliance to communicate with the firewall and Panorama running 8.0.1 and later.
28769 (5.1 and later) 28260 (5.0 and later) 49160 (5.0 and earlier)	TCP TCP TCP	Used for the HA connectivity and synchronization between Panorama HA peers using clear text communication. Communication can be initiated by either peer.
28	TCP	Used for the HA connectivity and synchronization between Panorama HA peers using encrypted communication (SSH over TCP). Communication can be initiated by either peer.
28270 (6.0 and later) 49190 (5.1 and earlier)	TCP	Used for communication among Log Collectors in a Collector Group for log distribution.
2049	TCP	Used by the Panorama virtual appliance to write logs to the NFS datastore.
23000 to 23999	TCP, UDP, or SSL	Used for Syslog communication between Panorama and the Traps ESM components.

Ports Used for GlobalProtect

GlobalProtect uses the following ports.

Destination Port	Protocol	Description
443	TCP	Used for communication between GlobalProtect agents and portals, or GlobalProtect agents and gateways and for SSL tunnel connections. GlobalProtect gateways also use this port to collect host information from GlobalProtect agents and perform host information profile (HIP) checks.
4501	UDP	Used for IPSec tunnel connections between GlobalProtect agents and gateways.

For tips on how to use a loopback interface to provide access to GlobalProtect on different ports and addresses, refer to Can GlobalProtect Portal Page be Configured to be Accessed on any Port?

© 2018 Palo Alto Networks, Inc.

Ports Used for User-ID

User-ID is a feature that enables mapping of user IP addresses to usernames and group memberships, enabling user- or group-based policy and visibility into user activity on your network (for example, to be able to quickly track down a user who may be the victim of a threat). To perform this mapping, the firewall, the User-ID agent (either installed on a Windows-based system or the PAN-OS integrated agent running on the firewall), and/or the Terminal Services agent must be able to connect to directory services on your network to perform Group Mapping and User Mapping. Additionally, if the agents are running on systems external to the firewall, they must be able to connect to the firewall to communicate the IP address to username mappings to the firewall. The following table lists the communication requirements for User-ID along with the port numbers required to establish connections.

Destination Port	Protocol	Description
389	TCP	Port the firewall uses to connect to an LDAP server (plaintext or Start Transport Layer Security (Start TLS) to Map Users to Groups.
3268	TCP	Port the firewall uses to connect to an Active Directory global catalog server (plaintext or Start TLS) to Map Users to Groups.
636	TCP	Port the firewall uses for LDAP over SSL connections with an LDAP server to Map Users to Groups.
3269	TCP	Port the firewall uses for LDAP over SSL connections with an Active Directory global catalog server to Map Users to Groups.
514 6514	TCP UDP SSL	Port the User-ID agent listens on for authentication syslog messages if you Configure User-ID to Monitor Syslog Senders for User Mapping. The port depends on the type of agent and protocol: • PAN-OS integrated User-ID agent—Port 6514 for SSL and port 514 for UDP. • Windows-based User-ID agent—Port 514 for both TCP and UDP.
5007	TCP	Port the firewall listens on for user mapping information from the User-ID or Terminal Services agent. The agent sends the IP address and username mapping along with a timestamp whenever it learns of a new or updated mapping. In addition, it connects to the firewall at regular intervals to refresh known mappings.
5006	TCP	Port the User-ID agent listens on for XML API requests. The source for this communication is typically the system running a script that invokes the API.
88	UDP/TCP	Port the User-ID agent uses to authenticate to a Kerberos server. The firewall tries UDP first and falls back to TCP.
1812	UDP	Port the User-ID agent uses to authenticate to a RADIUS server.
49	TCP	Port the User-ID agent uses to authenticate to a TACACS+ server.
135	TCP	Port the User-ID agent uses to establish TCP-based WMI connections with the Microsoft Remote Procedure Call (RPC) Endpoint Mapper. The

© 2018 Palo Alto Networks, Inc.

Destination Port	Protocol	Description
		Endpoint Mapper then assigns the agent a randomly assigned port in the 49152-65535 port range. The agent uses this connection to make RPC queries for Exchange Server or AD server security logs, session tables. This is also the port used to access Terminal Services. The User-ID agent also uses this port to connect to client systems to perform Windows Management Instrumentation (WMI) probing.
139	TCP	Port the User-ID agent uses to establish TCP-based NetBIOS connections to the AD server so that it can send RPC queries for security logs and session information. The User-ID agent also uses this port to connect to client systems for NetBIOS probing (supported on the Windows-based User-ID agent only).
445	TCP	Port the User-ID agent uses to connect to the Active Directory (AD) using TCP-based SMB connections to the AD server for access to user logon information (print spooler and Net Logon).

© 2018 Palo Alto Networks, Inc.

Reset the Firewall to Factory Default Settings

Resetting the firewall to factory defaults will result in the loss of all configuration settings and logs.

STEP 1 | Set up a console connection to the firewall.

1. Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1).

   If your computer does not have a 9-pin serial port, use a USB-to-serial port connector.

2. Enter your login credentials.
3. Enter the following CLI command:

   `debug system maintenance-mode`

   The firewall will reboot in the maintenance mode.

STEP 2 | Reset the system to factory default settings.

1. When the firewall reboots, press **`Enter`** to continue to the maintenance mode menu.
2. Select **`Factory Reset`** and press **`Enter`**.
3. Select **`Factory Reset`** and press **`Enter`** again.

   The firewall will reboot without any configuration settings. The default username and password to log in to the firewall is admin/admin.

   To perform initial configuration on the firewall and to set up network connectivity, see Integrate the Firewall into Your Management Network.

© 2018 Palo Alto Networks, Inc.

Bootstrap the Firewall

Bootstrapping speeds up the process of configuring and licensing the firewall to make it operational on the network with or without Internet access. Bootstrapping allows you to choose whether to configure the firewall with a basic configuration file (init-cfg.txt) so that it can connect to Panorama and obtain the complete configuration or to fully configure the firewall with the basic configuration and the optional bootstrap.xml file.

- USB Flash Drive Support
- Sample init-cfg.txt Files
- Prepare a USB Flash Drive for Bootstrapping a Firewall
- Bootstrap a Firewall Using a USB Flash Drive

USB Flash Drive Support

The USB flash drive that bootstraps a hardware-based Palo Alto Networks firewall must support one of the following:

- File Allocation Table 32 (FAT32)
- Third Extended File System (ext3)

The firewall can bootstrap from the following flash drives with USB2.0 or USB3.0 connectivity:

Supported USB Flash Drives
Kingston • Kingston SE9 8GB (2.0) • Kingston SE9 16GB (3.0) • Kingston SE9 32GB (3.0)
SanDisk • SanDisk Cruzer Fit CZ33 8GB (2.0) • SanDisk Cruzer Fit CZ33 16GB (2.0) • SanDisk Cruzer CZ36 16GB (2.0) • SanDisk Cruzer CZ36 32GB (2.0) • SanDisk Extreme CZ80 32GB (3.0)
Silicon Power • Silicon Power Jewel 32GB (3.0) • Silicon Power Blaze 16GB (3.0)
PNY • PNY Attache 16GB (2.0) • PNY Turbo 32GB (3.0)

© 2018 Palo Alto Networks, Inc.

Sample init-cfg.txt Files

An init-cfg.txt file is required for the bootstrap process; this file is a basic configuration file that you create using a text editor. To create this file, see Create the init-cfg.txt file. The following sample init-cfg.txt files show the parameters that are supported in the file; the parameters that you must provide are in bold.

Sample init-cfg.txt (Static IP Address)	Sample init-cfg.txt (DHCP Client)
type=static **ip-address=10.5.107.19** **default-gateway=10.5.107.1** **netmask=255.255.255.0** **ipv6-address=2001:400:f00::1/64** **ipv6-default-gateway=2001:400:f00::2** hostname=Ca-FW-DC1 panorama-server=10.5.107.20 panorama-server-2=10.5.107.21 tplname=FINANCE_TG4 dgname=finance_dg dns-primary=10.5.6.6 dns-secondary=10.5.6.7 op-command-modes=multi-vsys,jumbo-frame dhcp-send-hostname=no dhcp-send-client-id=no dhcp-accept-server-hostname=no dhcp-accept-server-domain=no	**type=dhcp-client** ip-address= default-gateway= netmask= ipv6-address= ipv6-default-gateway= hostname=Ca-FW-DC1 panorama-server=10.5.107.20 panorama-server-2=10.5.107.21 tplname=FINANCE_TG4 dgname=finance_dg dns-primary=10.5.6.6 dns-secondary=10.5.6.7 op-command-modes=multi-vsys,jumbo-frame dhcp-send-hostname=yes dhcp-send-client-id=yes dhcp-accept-server-hostname=yes dhcp-accept-server-domain=yes

The following table describes the fields in the init-cfg.txt file. The type is required; if the type is static, the IP address, default gateway and netmask are required, or the IPv6 address and IPv6 default gateway are required.

Field	Description
type	(Required) Type of management IP address: static or dhcp-client.
ip-address	(Required for IPv4 static management address) IPv4 address. The firewall ignores this field if the type is dhcp-client.
default-gateway	(Required for IPv4 static management address) IPv4 default gateway for the management interface. The firewall ignores this field if the type is dhcp-client.
netmask	(Required for IPv4 static management address) IPv4 netmask. The firewall ignores this field if the type is dhcp-client.
ipv6-address	(Required for IPv6 static management address) IPv6 address and /prefix length of the management interface. The firewall ignores this field if the type is dhcp-client.

© 2018 Palo Alto Networks, Inc.

Field	Description
ipv6-default-gateway	(Required for IPv6 static management address) IPv6 default gateway for the management interface. The firewall ignores this field if the type is dhcp-client.
hostname	(Optional) Host name for the firewall.
panorama-server	(Recommended) IPv4 or IPv6 address of the primary Panorama server.
panorama-server-2	(Optional) IPv4 or IPv6 address of the secondary Panorama server.
tplname	(Recommended) Panorama template name.
dgname	(Recommended) Panorama device group name.
dns-primary	(Optional) IPv4 or IPv6 address of the primary DNS server.
dns-secondary	(Optional) IPv4 or IPv6 address of the secondary DNS server.
vm-auth-key	(VM-Series firewalls only) Virtual machine authentication key.
op-command-modes	(Optional) Enter multi-vsys, jumbo-frame, or both separated by a comma only. Enables multiple virtual systems and jumbo frames while bootstrapping.
dhcp-send-hostname	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall sends its hostname to the DHCP server.
dhcp-send-client-id	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall sends its client ID to the DHCP server.
dhcp-accept-server-hostname	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall accepts its hostname from the DHCP server.
dhcp-accept-server-domain	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall accepts its DNS server from the DHCP server.

Prepare a USB Flash Drive for Bootstrapping a Firewall

You can use a USB flash drive to bootstrap a physical firewall. However, to do so you must upgrade to PAN-OS 8.0 and Reset the Firewall to Factory Default Settings. For security reasons, you can bootstrap a firewall only when it is in factory default state or has all private data deleted.

STEP 1 | Obtain serial numbers (S/Ns) and auth codes for support subscriptions from your order fulfillment email.

STEP 2 | Register S/Ns of new firewalls on the Customer Support portal.

1. Go to support.paloaltonetworks.com, log in, and select **Assets** > **Devices** > **Register New Device** > **Register device using Serial Number or Authorization Code**.
2. Follow the steps to Register the Firewall.
3. Click **Submit**.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Activate authorization codes on the Customer Support portal, which creates license keys.

1. Go to support.paloaltonetworks.com, log in, and select the **Assets** > **Devices** on the left-hand navigation pane.
2. For each device S/N you just registered, click the **Action** link (the pencil icon).
3. Under Activate Licenses, select **Activate Auth-Code**.
4. Enter the **Authorization code** and click **Agree** and **Submit**.

STEP 4 | Add the S/Ns in Panorama.

Complete Step 1 in Add a Firewall as a Managed Device in the Panorama Administrator's Guide.

STEP 5 | Create the init-cfg.txt file.

Create the init-cfg.txt file, a mandatory file that provides bootstrap parameters. The fields are described in Sample init-cfg.txt Files.



If the init-cfg.txt file is missing, the bootstrap process will fail and the firewall will boot up with the default configuration in the normal boot-up sequence.

There are no spaces between the key and value in each field; do not add spaces because they cause failures during parsing on the management server side.

You can have multiple init-cfg.txt files—one each for different remote sites—by prepending the S/N to the file name. For example:

0008C200105-init-cfg.txt

0008C200107-init-cfg.txt

If no prepended filename is present, the firewall uses the init-cfg.txt file and proceeds with bootstrapping.

STEP 6 | (Optional) Create the bootstrap.xml file.

The optional bootstrap.xml file is a complete firewall configuration that you can export from an existing production firewall.

1. Select **Device** > **Setup** > **Operations** > **Export named configuration snapshot**.
2. Select the **Name** of the saved or the running configuration.
3. Click **OK**.
4. Rename the file as **`bootstrap.xml`**.

STEP 7 | Create and download the bootstrap bundle from the Customer Support portal.

For a physical firewall, the bootstrap bundle requires only the /license and /config directories.

Use one of the following methods to create and download the bootstrap bundle:

- Use **Method 1** to create a bootstrap bundle specific to a remote site (you have only one init-cfg.txt file).
- Use **Method 2** to create one bootstrap bundle for multiple sites.

Method 1

1. On your local system, go to support.paloaltonetworks.com and log in.
2. Select **Assets**.
3. Select the S/N of the firewall you want to bootstrap.
4. Select **Bootstrap Container**.
5. Click **Select**.

© 2018 Palo Alto Networks, Inc.

6. Upload and **Open** the init-cfg.txt file you created.
7. (Optional) Select the bootstrap.xml file you created and **Upload Files**.

   You must use a bootstrap.xml file from a firewall of the same model and PAN-OS version.
8. Select **Bootstrap Container Download** to download a tar.gz file named **bootstrap_<S/N>_<date>.tar.gz** to your local system. This bootstrap container includes the license keys associated with the S/N of the firewall.

Method 2

Create a tar.gz file on your local system with two top-level directories: /license and /config. Include all licenses and all init-cfg.txt files with S/Ns prepended to the filenames.

The license key files you download from the Customer Support portal have the S/N in the license file name. PAN-OS checks the S/N in the file name against the firewall S/N while executing the bootstrap process.

STEP 8 | Import the tar.gz file you created to a PAN-OS 7.1 firewall using Secure Copy (SCP) or TFTP.

Access the CLI and enter one of the following commands:

- **tftp import bootstrap-bundle file <path and filename> from <host IP address>**

  For example:

  tftp import bootstrap-bundle file /home/userx/bootstrap/devices/pa5000.tar.gz from 10.1.2.3
- **scp import bootstrap-bundle from <<user>@<host>:<path to file>>**

  For example:

  scp import bootstrap-bundle from userx@10.1.2.3:/home/userx/bootstrap/devices/pa200_bootstrap_bundle.tar.gz

STEP 9 | Prepare the USB flash drive.

1. Insert the USB flash drive into the firewall that you used in Step 8.
2. Enter the following CLI operational command, using your tar.gz filename in place of "**pa5000.tar.gz**". This command formats the USB flash drive, unzips the file, and validates the USB flash drive:

   request system bootstrap-usb prepare from pa5000.tar.gz
3. Press **y** to continue. The following message displays when the USB drive is ready:

   ```
   USB prepare completed successfully.
   ```
4. Remove the USB flash drive from the firewall.
5. You can prepare as many USB flash drives as needed.

STEP 10 | Deliver the USB flash drive to your remote site.

If you used Method 2 to create the bootstrap bundle, you can use the same USB flash drive content for bootstrapping firewalls at multiple remote sites. You can translate the content into multiple USB flash drives or a single USB flash drive used multiple times.

Bootstrap a Firewall Using a USB Flash Drive

After you receive a new Palo Alto Networks firewall and a USB flash drive loaded with bootstrap files, you can bootstrap the firewall.

© 2018 Palo Alto Networks, Inc.

Microsoft Windows and Apple Mac operating systems are unable to read the bootstrap USB flash drive because the drive is formatted using an ext4 file system. You must install third-party software or use a Linux system to read the USB drive.

STEP 1 | The firewall must be in a factory default state or must have all private data deleted.

STEP 2 | To ensure connectivity with your corporate headquarters, cable the firewall by connecting the management interface (MGT) using an Ethernet cable to one of the following:

- An upstream modem
- A port on the switch or router
- An Ethernet jack in the wall

STEP 3 | Insert the USB flash drive into the USB port on the firewall and power on the firewall. The factory default firewall bootstraps itself from the USB flash drive.

The firewall Status light turns from yellow to green when the firewall is configured; autocommit is successful.

STEP 4 | Verify bootstrap completion. You can see basic status logs on the console during the bootstrap and you can verify that the process is complete.

1. If you included Panorama values (panorama-server, tplname, and dgname) in your init-cfg.txt file, check Panorama managed devices, device group, and template name.
2. Verify the general system settings and configuration by accessing the web interface and selecting **Dashboard** > **Widgets** > **System** or by using the CLI operational commands **`show system info`** and **`show config running`**.
3. Verify the license installation by selecting **Device** > **Licenses** or by using the CLI operational command **`request license info`**.
4. If you have Panorama configured, manage the content versions and software versions from Panorama. If you do not have Panorama configured, use the web interface to manage content versions and software versions.

© 2018 Palo Alto Networks, Inc.

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. You can choose from several authentication services to protect your network and to accommodate your existing security infrastructure while ensuring a smooth user experience.

If you have a public key infrastructure, you can deploy certificates to enable authentication without users having to manually respond to login challenges (see Certificate Management). Alternatively, or in addition to certificates, you can implement interactive authentication, which requires users to authenticate using one or more methods. The following topics describe how to implement, test, and troubleshoot the different types of interactive authentication:

> Authentication Types
> Plan Your Authentication Deployment
> Configure Multi-Factor Authentication
> Configure SAML Authentication
> Configure Kerberos Single Sign-On
> Configure Kerberos Server Authentication
> Configure TACACS+ Authentication
> Configure RADIUS Authentication
> Configure LDAP Authentication
> Connection Timeouts for Authentication Servers
> Configure Local Database Authentication
> Configure an Authentication Profile and Sequence
> Test Authentication Server Connectivity
> Authentication Policy
> Troubleshoot Authentication Issues

© 2018 Palo Alto Networks, Inc.

Authentication Types

- External Authentication Services
- Multi-Factor Authentication
- SAML
- Kerberos
- TACACS+
- RADIUS
- LDAP
- Local Authentication

External Authentication Services

The firewall and Panorama can use external servers to control administrative access to the web interface and end user access to services or applications through Captive Portal and GlobalProtect. In this context, any authentication service that is not local to the firewall or Panorama is considered external, regardless of whether the service is internal (such as Kerberos) or external (such as a SAML identity provider) relative to your network. The server types that the firewall and Panorama can integrate with include Multi-Factor Authentication (MFA), SAML, Kerberos, TACACS+, RADIUS, and LDAP. Although you can also use the Local Authentication services that the firewall and Panorama support, usually external services are preferable because they provide:

- Central management of all user accounts in an external identity store. All the supported external services provide this option for end users and administrators.
- Central management of account authorization (role and access domain assignments). SAML, TACACS+, and RADIUS support this option for administrators.
- Single sign-on (SSO), which enables users to authenticate only once for access to multiple services and applications. SAML and Kerberos support SSO.
- Multiple authentication challenges of different types (factors) to protect your most sensitive services and applications. MFA services support this option.

Authentication through an external service requires a server profile that defines how the firewall connects to the service. You assign the server profile to authentication profiles, which define settings that you customize for each application and set of users. For example, you can configure one authentication profile for administrators who access the web interface and another profile for end users who access a GlobalProtect portal. For details, see Configure an Authentication Profile and Sequence.

Multi-Factor Authentication

You can Configure Multi-Factor Authentication (MFA) to ensure that each user authenticates using multiple methods (factors) when accessing highly sensitive services and applications. For example, you can force users to enter a login password and then enter a verification code that they receive by phone before allowing access to important financial documents. This approach helps to prevent attackers from accessing every service and application in your network just by stealing passwords. Of course, not every service and application requires the same degree of protection, and MFA might not be necessary for less sensitive services and applications that users access frequently. To accommodate a variety of security needs, you can Configure Authentication Policy rules that trigger MFA or a single authentication factor (such as login credentials or certificates) based on specific services, applications, and end users.

When choosing how many and which types of authentication factors to enforce, it's important to understand how policy evaluation affects the user experience. When a user requests a service or application, the firewall first evaluates Authentication policy. If the request matches an Authentication policy rule with MFA enabled, the firewall displays a Captive Portal web form so that users can authenticate

© 2018 Palo Alto Networks, Inc.

for the first factor. If authentication succeeds, the firewall displays an MFA login page for each additional factor. Some MFA services prompt the user to choose one factor out of two to four, which is useful when some factors are unavailable. If authentication succeeds for all factors, the firewall evaluates Security policy for the requested service or application.



To reduce the frequency of authentication challenges that interrupt the user workflow, you can configure the first factor to use Kerberos *or* SAML *single sign-on (SSO) but not* NT LAN Manager (NTLM) *authentication.*

To implement MFA for GlobalProtect, refer to Configure GlobalProtect to Facilitate Multi-FactorAuthenticationNotifications.

You cannot use MFA authentication profiles in authentication sequences.

For end-user authentication via Authentication Policy, the firewall directly integrates with several MFA platforms (Duo v2, Okta Adaptive, PingID, and RSA SecurID), as well as integrating through RADIUS or SAML for all other MFA platforms. For remote user authentication to GlobalProtect portals and gateways and for administrator authentication to the Panorama and PAN-OS web interface, the firewall integrates with MFA vendors using RADIUS and SAML only.

The firewall supports the following MFA factors:

Factor	Description
Push	An endpoint device (such as a phone or tablet) prompts the user to allow or deny authentication.
Short message service (SMS)	An SMS message on the endpoint device prompts the user to allow or deny authentication. In some cases, the endpoint device provides a code that the user must enter in the MFA login page.
Voice	An automated phone call prompts the user to authenticate by pressing a key on the phone or entering a code in the MFA login page.
One-time password (OTP)	An endpoint device provides an automatically generated alphanumeric string, which the user enters in the MFA login page to enable authentication for a single transaction or session.

SAML

You can use Security Assertion Markup Language (SAML) 2.0 to authenticate administrators who access the firewall or Panorama web interface and end users who access web applications that are internal or external to your organization. In environments where each user accesses many applications and authenticating for each one would impede user productivity, you can configure SAML single sign-on (SSO) to enable one login to access multiple applications. Likewise, SAML single logout (SLO) enables a user to end sessions for multiple applications by logging out of just one session. SSO is available to administrators who access the web interface and to end users who access applications through GlobalProtect or Captive Portal. SLO is available to administrators and GlobalProtect end users, but not to Captive Portal end users. When you configure SAML authentication on the firewall or on Panorama, you can specify SAML attributes for administrator authorization. SAML attributes enable you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on the firewall or Panorama.

© 2018 Palo Alto Networks, Inc.

Administrators cannot use SAML to authenticate to the firewall or Panorama CLI.

You cannot use SAML authentication profiles in authentication sequences.

SAML authentication requires a *service provider* (the firewall or Panorama), which controls access to applications, and an *identity provider* (IdP) such as PingFederate, which authenticates users. When a user requests a service or application, the firewall or Panorama intercepts the request and redirects the user to the IdP for authentication. The IdP then authenticates the user and returns a *SAML assertion*, which indicates authentication succeeded or failed. SAML Authentication for Captive Portal End Users illustrates SAML authentication for an end user who accesses applications through Captive Portal.

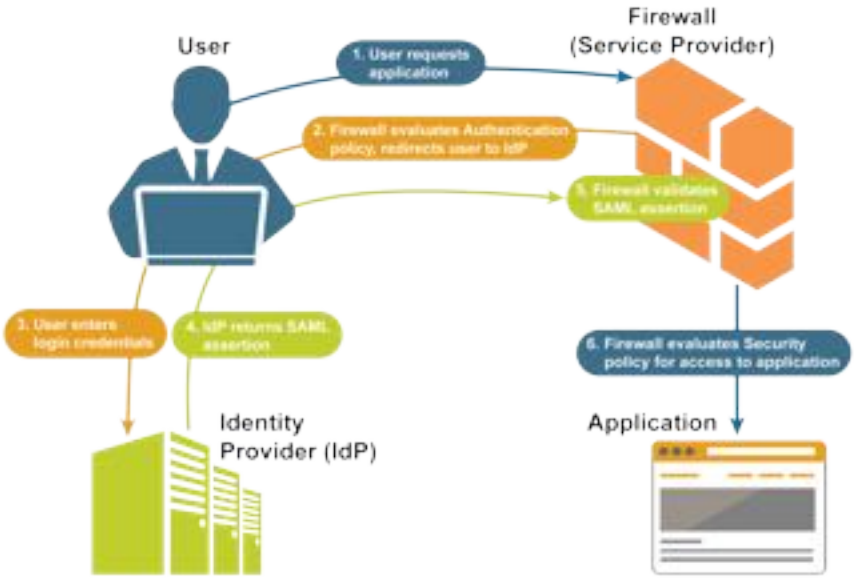


Figure 1: SAML Authentication for Captive Portal End Users

Kerberos

Kerberos is an authentication protocol that enables a secure exchange of information between parties over an insecure network using unique keys (called tickets) to identify the parties. The firewall and Panorama support two types of Kerberos authentication for administrators and end users:

- **Kerberos server authentication**—A Kerberos server profile enables users to natively authenticate to an Active Directory domain controller or a Kerberos V5-compliant authentication server. This authentication method is interactive, requiring users to enter usernames and passwords. For the configuration steps, see Configure Kerberos Server Authentication.
- **Kerberos single sign-on (SSO)**—A network that supports Kerberos V5 SSO prompts a user to log in only for initial access to the network (such as logging in to Microsoft Windows). After this initial login, the user can access any browser-based service in the network (such as the firewall web interface) without having to log in again until the SSO session expires. (Your Kerberos administrator sets the duration of SSO sessions.) If you enable both Kerberos SSO and another external authentication service (such as a TACACS+ server), the firewall first tries SSO and, only if that fails, falls back to the external service for authentication. To support Kerberos SSO, your network requires:
  - A Kerberos infrastructure, including a key distribution center (KDC) with an authentication server (AS) and ticket-granting service (TGS).
  - A Kerberos account for the firewall or Panorama that will authenticate users. An account is required to create a Kerberos keytab, which is a file that contains the principal name and hashed password of the firewall or Panorama. The SSO process requires the keytab.

  For the configuration steps, see Configure Kerberos Single Sign-On.

© 2018 Palo Alto Networks, Inc.

Kerberos SSO is available only for services and applications that are internal to your Kerberos environment. To enable SSO for external services and applications, use SAML.

TACACS+

Terminal Access Controller Access-Control System Plus (TACACS+) is a family of protocols that enable authentication and authorization through a centralized server. TACACS+ encrypts usernames and passwords, making it more secure than RADIUS, which encrypts only passwords. TACACS+ is also more reliable because it uses TCP, whereas RADIUS uses UDP. You can configure TACACS+ authentication for end users or administrators on the firewall and for administrators on Panorama. Optionally, you can use TACACS+ Vendor-Specific Attributes (VSAs) to manage administrator authorization. TACACS+ VSAs enable you to quickly change the roles, access domains, and user groups of administrators through your directory service instead of reconfiguring settings on the firewall and Panorama.

The firewall and Panorama support the following TACACS+ attributes and VSAs. Refer to your TACACS+ server documentation for the steps to define these VSAs on the TACACS+ server.

Name	Value
service	This attribute is required to identify the VSAs as specific to Palo Alto Networks. You must set the value to **`paloalto`**.
protocol	This attribute is required to identify the VSAs as specific to Palo Alto Networks devices. You must set the value to **`firewall`**.
PaloAlto-Admin-Role	A default (dynamic) administrative role name or a custom administrative role name on the firewall.
PaloAlto-Admin-Access-Domain	The name of an access domain for firewall administrators (configured in the **Device** > **Access Domains** page). Define this VSA if the firewall has multiple virtual systems.
PaloAlto-Panorama-Admin-Role	A default (dynamic) administrative role name or a custom administrative role name on Panorama.
PaloAlto-Panorama-Admin-Access-Domain	The name of an access domain for Device Group and Template administrators (configured in the **Panorama** > **Access Domains** page).
PaloAlto-User-Group	The name of a user group in the Allow List of an authentication profile.

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is a broadly supported networking protocol that provides centralized authentication and authorization. You can configure RADIUS authentication for end users or administrators on the firewall and for administrators on Panorama. Optionally, you can use RADIUS Vendor-Specific Attributes (VSAs) to manage administrator authorization. RADIUS VSAs enable you to quickly change the roles, access domains, and user groups of administrators through your directory service instead of reconfiguring settings on the firewall and Panorama. You can also configure the firewall to use a RADIUS server for:

© 2018 Palo Alto Networks, Inc.

- Collecting VSAs from GlobalProtect clients.
- Implementing Multi-Factor Authentication.

When sending authentication requests to a RADIUS server, the firewall and Panorama use the authentication profile name as the network access server (NAS) identifier, even if the profile is assigned to an authentication sequence for the service (such as administrative access to the web interface) that initiates the authentication process.

The firewall and Panorama support the following RADIUS VSAs. To define VSAs on a RADIUS server, you must specify the vendor code (25461 for Palo Alto Networks firewalls or Panorama) and the VSA name and number. Some VSAs also require a value. Refer to your RADIUS server documentation for the steps to define these VSAs.

When configuring the advanced vendor options on a Cisco Secure Access Control Server (ACS), you must set both the Vendor Length Field Size and Vendor Type Field Size to ***1****. Otherwise, authentication will fail.*

Name	Number	Value
VSAs for administrator account management andauthentication		
PaloAlto-Admin-Role	1	A default (dynamic) administrative role name or a custom administrative role name on the firewall.
PaloAlto-Admin-Access-Domain	2	The name of an access domain for firewall administrators (configured in the **Device** > **Access Domains** page). Define this VSA if the firewall has multiple virtual systems.
PaloAlto-Panorama-Admin-Role	3	A default (dynamic) administrative role name or a custom administrative role name on Panorama.
PaloAlto-Panorama-Admin-Access-Domain	4	The name of an access domain for Device Group and Template administrators (configured in the **Panorama** > **Access Domains** page).
PaloAlto-User-Group	5	The name of a user group that an authentication profile references.
VSAs forwarded from GlobalProtect clients to the RADIUSserver		
PaloAlto-User-Domain	6	Don’t specify a value when you define these VSAs.
PaloAlto-Client-Source-IP	7	
PaloAlto-Client-OS	8	
PaloAlto-Client-Hostname	9	
PaloAlto-GlobalProtect-Client-Version	10	

© 2018 Palo Alto Networks, Inc.

LDAP

Lightweight Directory Access Protocol (LDAP) is a standard protocol for accessing information directories. You can Configure LDAP Authentication for end users and for firewall and Panorama administrators.

Configuring the firewall to connect to an LDAP server also enables you to define policy rules based on users and user groups instead of just IP addresses. For the steps, see Map Users to Groups and Enable User- and Group-Based Policy.

Local Authentication

Although the firewall and Panorama provide local authentication for administrators and end users, External Authentication Services are preferable in most cases because they provide central account management. However, you might require special user accounts that you don't manage through the directory servers that your organization reserves for regular accounts. For example, you might define a superuser account that is local to the firewall so that you can access the firewall even if the directory server is down. In such cases, you can use the following local authentication methods:

- **(Firewall only) Local database authentication**—To Configure Local Database Authentication, you create a database that runs locally on the firewall and contains user accounts (usernames and passwords or hashed passwords) and user groups. This type of authentication is useful for creating user accounts that reuse the credentials of existing Unix accounts in cases where you know only the hashed passwords, not the plaintext passwords. Because local database authentication is associated with authentication profiles, you can accommodate deployments where different sets of users require different authentication settings, such as Kerberos single sign-on (SSO) or Multi-Factor Authentication (MFA). (For details, see Configure an Authentication Profile and Sequence). For accounts that use plaintext passwords, you can also define password complexity and expiration settings. This authentication method is available to administrators who access the firewall (but not Panorama) and end users who access services and applications through Captive Portal or GlobalProtect.
- **Local authentication without a database**—You can configure firewall administrative accounts or Panorama administrative accounts without creating a database of users and user groups that runs locally on the firewall or Panorama. Because this method is not associated with authentication profiles, you cannot combine it with Kerberos SSO or MFA. However, this is the only authentication method that allows password profiles, which enable you to associate individual accounts with password expiration settings that differ from the global settings. (For details, see Define password complexity and expiration settings)

© 2018 Palo Alto Networks, Inc.

Plan Your Authentication Deployment

The following are key questions to consider before you implement an authentication solution for administrators who access the firewall and end users who access services and applications through Captive Portal.

For both end users and administrators, consider:

- How can you leverage your existing security infrastructure? Usually, integrating the firewall with an existing infrastructure is faster and cheaper than setting up a new, separate solution just for firewall services. The firewall can integrate with Multi-Factor Authentication, SAML, Kerberos, TACACS+, RADIUS, and LDAP servers. If your users access services and applications that are external to your network, you can use SAML to integrate the firewall with an identity provider (IdP) that controls access to both external and internal services and applications.
- How can you optimize the user experience? If you don't want users to authenticate manually and you have a public key infrastructure, you can implement certificate authentication. Another option is to implement Kerberos or SAML single sign-on (SSO) so that users can access multiple services and applications after logging in to just one. If your network requires additional security, you can combine certificate authentication with interactive (challenge-response) authentication.
- Do you require special user accounts that you don't manage through the directory servers that your organization reserves for regular accounts? For example, you might define a superuser account that is local to the firewall so that you can access the firewall even if the directory server is down. You can configure Local Authentication for these special-purpose accounts.



External Authentication Services *are usually preferable to local authentication because they provide central account management.*

For end users only, consider:

- Which services and applications are more sensitive than others? For example, you might want stronger authentication for key financial documents than for search engines. To protect your most sensitive services and applications, you can configure Multi-Factor Authentication (MFA) to ensure that each user authenticates using multiple methods (factors) when accessing those services and applications. To accommodate a variety of security needs, Configure Authentication Policy rules that trigger MFA or single factor authentication (such as login credentials or certificates) based on specific services, applications, and end users. Other ways to reduce your attack surface include network segmentation and user groups for whitelist applications.

For administrators only, consider:

- Do you use an external server to centrally manage authorization for all administrative accounts? By defining Vendor-Specific Attributes (VSAs) on the external server, you can quickly change administrative role assignments through your directory service instead of reconfiguring settings on the firewall. VSAs also enable you to specify access domains for administrators of firewalls with multiple virtual systems. SAML, TACACS+, and RADIUS support external authorization.

© 2018 Palo Alto Networks, Inc.

Configure Multi-Factor Authentication

To use Multi-Factor Authentication (MFA) for protecting sensitive services and applications, you must configure Captive Portal to display a web form for the first authentication factor and to record Authentication Timestamps. The firewall uses the timestamps to evaluate the timeouts for Authentication Policy rules. To enable additional authentication factors, you can integrate the firewall with MFA vendors through RADIUS or vendor APIs. After evaluating Authentication policy, the firewall evaluates Security policy, so you must configure rules for both policy types.



Palo Alto Networks provides support for MFA vendor API integrations through Applications or Applications and Threats content updates. This means that if you use Panorama to push device group configurations to firewalls, you must install the same Applications updates on the firewalls as on Panorama to avoid mismatches in vendor support.

MFA vendor API integrations are supported for end-user authentication through Authentication Policy only. For remote user authentication to GlobalProtect portals or gateways or for administrator authentication to the PAN-OS or Panorama web interface, you can only use MFA vendors supported through RADIUS or SAML; MFA services through vendor APIs are not supported in these use cases.

STEP 1 | Configure Captive Portal in **Redirect** mode to display a web form for the first authentication factor, to record authentication timestamps, and to update user mappings.

STEP 2 | Configure one of the following server profiles to define how the firewall will connect to the service that authenticates users for the first authentication factor.

- Add a RADIUS server profile. This is required if the firewall integrates with an MFA vendor through RADIUS. In this case, the MFA vendor provides the first and all additional authentication factors, so you can skip the next step (configuring an MFA server profile). If the firewall integrates with an MFA vendor through an API, you can still use a RADIUS server profile for the first factor but MFA server profiles are required for the additional factors.
- Add a SAML IdP server profile.
- Add a Kerberos server profile.
- Add a TACACS+ server profile.
- Add an LDAP server profile.



In most cases, an external service is recommended for the first authentication factor. However, you can configure Configure Local Database Authentication as an alternative.

STEP 3 | Add an MFA server profile.

The profile defines how the firewall connects to the MFA server. Add a separate profile for each authentication factor after the first factor. The firewall integrates with these MFA servers through vendor APIs. You can specify up to three additional factors. Each MFA vendor provides one factor, though some vendors let users choose one factor out of several.

1. Select **Device** > **Server Profiles** > **Multi Factor Authentication** and **Add** a profile.
2. Enter a **Name** to identify the MFA server.
3. Select the **Certificate Profile** that the firewall will use to validate the MFA server certificate when establishing a secure connection to the MFA server.
4. Select the **MFA Vendor** you deployed.
5. Configure the **Value** of each vendor attribute.

© 2018 Palo Alto Networks, Inc.

The attributes define how the firewall connects to the MFA server. Each vendor **Type** requires different attributes and values; refer to your vendor documentation for details.

6. Click **OK** to save the profile.

STEP 4 | Configure an authentication profile.

The profile defines the order of the authentication factors that users must respond to.

1. Select **Device** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the authentication profile.
3. Select the **Type** for the first authentication factor and select the corresponding **Server Profile**.
4. Select **Factors**, **Enable Additional Authentication Factors**, and **Add** the MFA server profiles you configured.

   The firewall will invoke each MFA service in the listed order, from top to bottom.
5. Click **OK** to save the authentication profile.

STEP 5 | 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 the **Authentication Profile** you configured and enter a **Message** that tells users how to authenticate for the first factor. The message displays in the Captive Portal web form.



If you set the Authentication Method to browser-challenge, the Captive Portal web form displays only if Kerberos SSO authentication fails. Otherwise, authentication for the first factor is automatic; users won't see the web form.

STEP 6 | Configure an Authentication policy rule.

The rule must match the services and applications you want to protect and the users who must authenticate.

1. Select **Policies** > **Authentication** and **Add** a rule.
2. Enter a **Name** to identify the rule.
3. Select **Source** and **Add** specific zones and IP addresses or select **Any** zones or IP addresses.

   The rule applies only to traffic coming from the specified IP addresses or from interfaces in the specified zones.
4. Select **User** and select or **Add** the source users and user groups to which the rule applies (default is **any**).
5. Select **Destination** and **Add** specific zones and IP addresses or select **any** zones or IP addresses.

   The IP addresses can be resources (such as servers) for which you want to control access.
6. Select **Service/URL Category** and select or **Add** the services and service groups for which the rule controls access (default is **service-http**).
7. 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 your most sensitive internal sites.
8. Select **Actions** and select the **Authentication Enforcement** object you created.
9. Specify the **Timeout** period in minutes (default 60) during which the firewall prompts the user to authenticate only once for repeated access to services and applications.
10. Click **OK** to save the rule.

STEP 7 | Customize the MFA login page.

© 2018 Palo Alto Networks, Inc.

The firewall displays this page to tell users how to authenticate for MFA factors and to indicate the authentication status (in progress, succeeded, or failed).

1. Select **Device** > **Response Pages** and select **MFA Login Page**.
2. Select the **Predefined** response page and **Export** the page to your client system.
3. On your client system, use an HTML editor to customize the downloaded response page and save it with a unique filename.
4. Return to the MFA Login Page dialog on the firewall, **Import** your customized page, **Browse** to select the **Import File**, select the **Destination** (virtual system or **shared** location), click **OK**, and click **Close**.

STEP 8 | Configure a Security policy rule that allows users to access the services and applications that require authentication.

1. Create a Security Policy Rule.
2. **Commit** your changes.



The automated correlation engine on the firewall uses several correlation objects to detect events on your network that could indicate credential abuse relating to MFA. To review the events, select Monitor > Automated Correlation Engine > Correlated Events.

STEP 9 | Verify that the firewall enforces MFA.

1. Log in to your network as one of the source users specified in the Authentication rule.
2. Request a service or application that matches one of the services or applications specified in the rule.

   The firewall displays the Captive Portal web form for the first authentication factor. The page contains the message you entered in the authentication enforcement object. For example:



3. Enter your user credentials for the first authentication challenge.

   The firewall then displays an MFA login page for the next authentication factor. For example, the MFA service might prompt you to select the Voice, SMS, push, or PIN code (OTP) authentication method. If you select push, your phone prompts you to approve the authentication.



4. Authenticate for the next factor.

   The firewall displays an authentication success or failure message. If authentication succeeded, the firewall displays an MFA login page for the next authentication factor, if any.

   Repeat this step for each MFA factor. After you authenticate for all the factors, the firewall evaluates Security policy to determine whether to allow access to the service or application.

© 2018 Palo Alto Networks, Inc.

5. End the session for the service or application you just accessed.
6. Start a new session for the same service or application. Be sure to perform this step within the **Timeout** period you configured in the Authentication rule.

   The firewall allows access without re-authenticating.
7. Wait until the **Timeout** period expires and request the same service or application.

   The firewall prompts you to re-authenticate.

Configure MFA Between RSA SecurID and the Firewall

Multi-factor authentication allows you to protect company assets by using multiple factors to verify a user's identity before allowing them to access network resources. To enable multi-factor authentication (MFA) between the firewall and the RSA SecurID Access Cloud Authentication Service, you must first configure the RSA SecurID Service so that you have the details that you need to configure the firewall to authenticate users using multiple factors. After you have performed the required configuration on the RSA SecurID Access Console, you can configure the firewall to integrate with RSA SecurID.

The Palo Alto Networks next-generation firewall integrates with the RSA SecurID Access Cloud Authentication Service. The MFA API integration with RSA SecurID is supported for cloud-based services only and does not support two-factor authentication for the on-premise Authentication Manager when the second factor uses the Vendor Specific API. The minimum content version required for this integration is 752 and PAN-OS 8.0.2.

- Get the RSA SecurID Access Cloud Authentication Service Details
- Configure the Firewall for MFA with RSA SecurID

Get the RSA SecurID Access Cloud Authentication Service Details

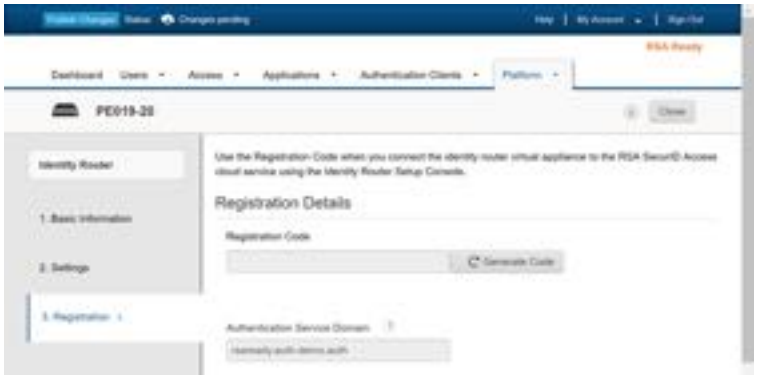
In order to securely pass user authentication requests to and from the firewall and the RSA SecurID Access Cloud Authentication Service, you must first go to the RSA SecurID Access Console and configure the RSA Access ID, the authentication service URL, and the client API key that the firewall needs to authenticate to and interact with the service. The firewall also needs the Access Policy ID that uses either the RSA Approve or RSA Tokencode authentication method to authenticate to the identity source.

- **Generate the RSA SecurID API key**—Log on to RSA SecurID Access Console and select **My Account** > **Company Settings** > **Authentication API Keys**. **Add** a new key and then **Save Settings** and **Publish Changes**.

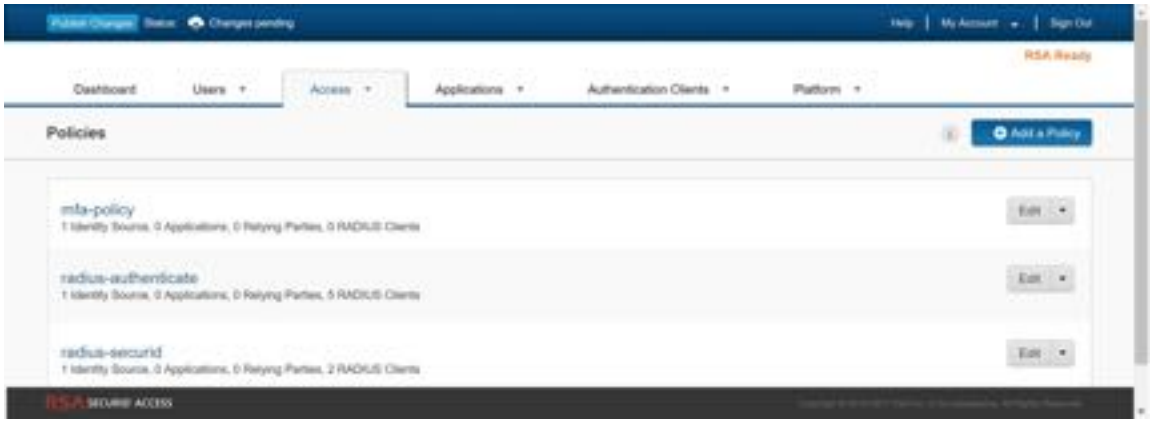


- **Get the RSA SecurID Access endpoint API (Authentication Service Domain) to which the firewall must connect**—Select **Platform** > **Identity Routers**, pick an Identity Router to **Edit** and jot down the **Authentication Service Domain**. In this example it is https://rsaready.auth-demo.auth.

© 2018 Palo Alto Networks, Inc.



- **Get the Access Policy ID**—Select **Access** > **Policies** and jot down the name of the access policy that will allow the firewall to act as an authentication client to the RSA SecurID service. The policy must be configured to use either the RSA Approve or the RSA Tokencode authentication methods only.



Configure the Firewall for MFA with RSA SecurID

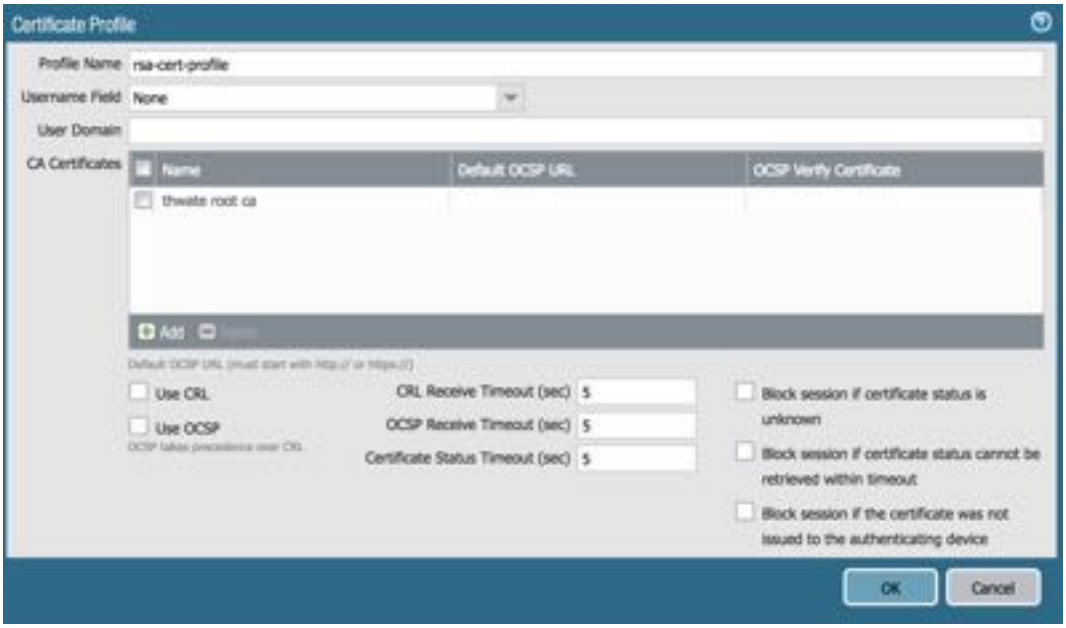
After you Get the RSA SecurID Access Cloud Authentication Service Details, you can configure the firewall to prompt users for an RSA SecurID token when MFA is invoked.

STEP 1 | Configure the firewall to trust the SSL certificate provided by the RSA SecurID Access endpoint API.

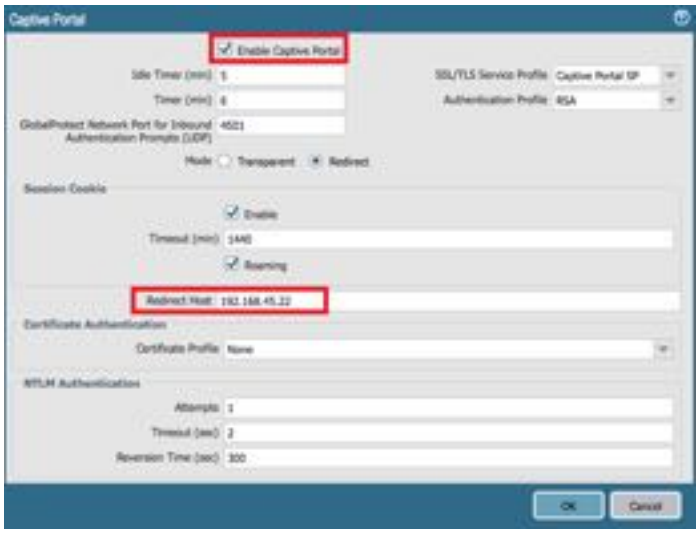
1. Export the SSL certificate from the RSA SecurID Access endpoint and import it into the firewall.

   To enable trust between the firewall and the RSA SecurID Access endpoint API, you must either import a self-signed certificate, or the CA certificate used to sign the certificate.
2. Configure a Certificate Profile (**Device** > **Certificate Management** > **Certificate Profile** and click **Add**).

© 2018 Palo Alto Networks, Inc.



STEP 2 | Configure Captive Portal (**Device** > **User Identification** > **Captive Portal Settings**) in Redirect mode to display a web form for authenticating to RSA SecureID. Make sure to specify the Redirect Host as an IP address or a hostname (with no periods in its name) that resolves to the IP address of the Layer 3 interface on the firewall to which web requests are redirected.



STEP 3 | Configure a multi-factor authentication server profile to specify how the firewall must connect with the RSA SecurID cloud service **(Device** > **Server Profiles** > **Multi Factor Authentication** and click **Add**).

1. Enter a **Name** to identify the MFA server profile.
2. Select the **Certificate Profile** that you created earlier, rsa-cert-profile in this example. The firewall will use this certificate when establishing a secure connection with RSA SecurID cloud service.
3. In the **MFA Vendor** drop-down, select **RSA SecurID Access**.
4. Configure the **Value** for each attribute that you noted in Get the RSA SecurID Access Cloud Authentication Service Details:

© 2018 Palo Alto Networks, Inc.

- **API Host**—Enter the hostname or IP address of the RSA SecurID Access API endpoint to which the firewall must connect, rsaready.auth-demo.auth in this example.
- **Base URI** —Do not modify the default value (/mfa/v1_1)
- **Client Key**—Enter the RSA SecurID Client Key.
- **Access ID**—Enter the RSA SecurID Access ID.
- **Assurance Policy**—Enter the RSA SecurID Access Policy name, mfa-policy in this example.
- **Timeout**—The default is 30 seconds.



5. Save the profile.

STEP 4 | Configure an authentication profile (**Device** > **Authentication Profile** and click **Add**).

The profile defines the order of the authentication factors that users must respond to.

1. Select the **Type** for the first authentication factor and select the corresponding **Server Profile**.
2. Select **Factors**, **Enable Additional Authentication Factors**, and **Add** the rsa-mfa server profile you created earlier in this example.



3. Click **OK** to save the authentication profile.

© 2018 Palo Alto Networks, Inc.

STEP 5 | Configure an authentication enforcement object. (**Objects** > **Authentication** and click **Add**).

Make sure to select the authentication profile you just defined called RSA in this example.



STEP 6 | Configure an Authentication policy rule. (**Policies** > **Authentication** and click **Add**)

Your authentication policy rule must match the services and applications you want to protect, specify the users who must authenticate, and include the authentication enforcement object that triggers the authentication profile. In this example, RSA SecurID Access authenticates all users who accessing HTTP, HTTPS, SSH, and VNC traffic with the authentication enforcement object called RSA Auth Enforcement (in **Actions**, select the **Authentication Enforcement** object).

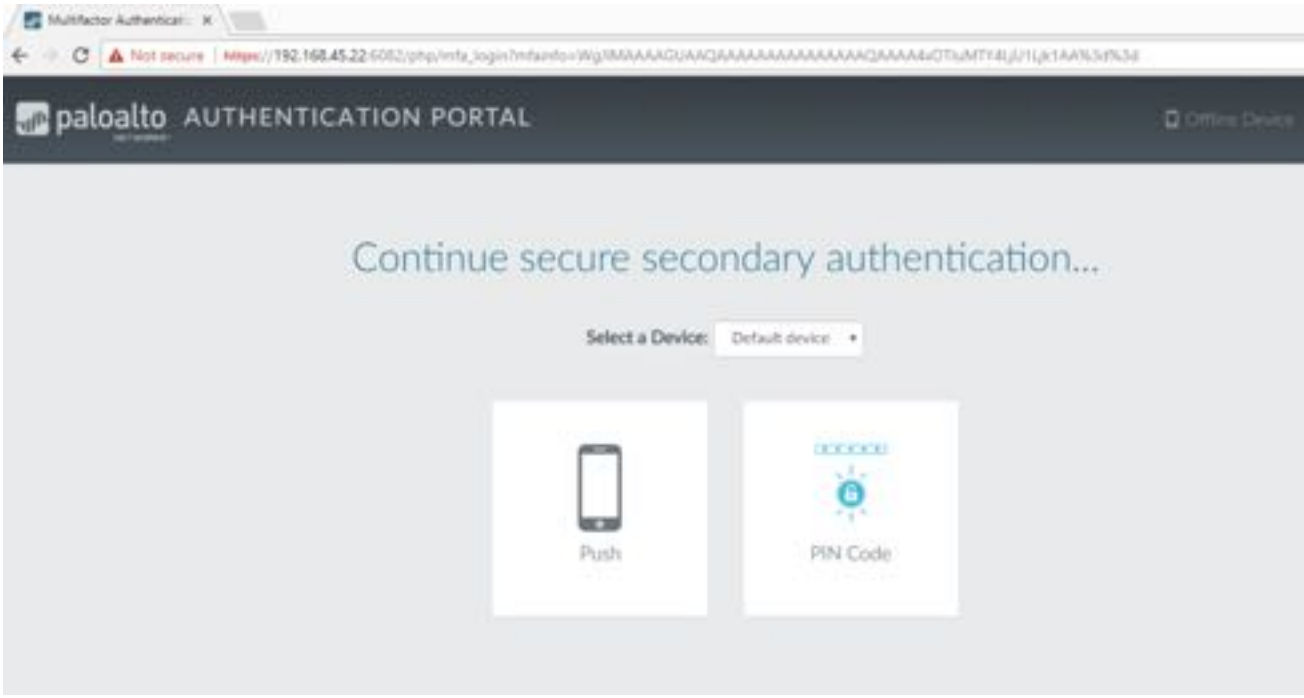


STEP 7 | **Commit** your changes on the firewall.

STEP 8 | Verify that users on your network are being secured using RSA SecurID using the Push or PIN Code authentication method you enabled.

1. **Push authentication**
    1. Ask a user on your network to launch a web browser and access a website. The Captive Portal page with the IP address or hostname for the Redirect Host you defined earlier should display.
    2. Verify that the user enters the credentials for the first authentication factor and then continues to the secondary authentication factor, and selects **Push**.

© 2018 Palo Alto Networks, Inc.



3. Check for a **Sign-In request** on the RSA SecurID Access application on the user's mobile device.
4. Ask the user to **Accept** the Sign-In Request on the mobile device, and wait for a few seconds for the firewall to receive the notification of successful authentication. The user should be able to access the requested website.

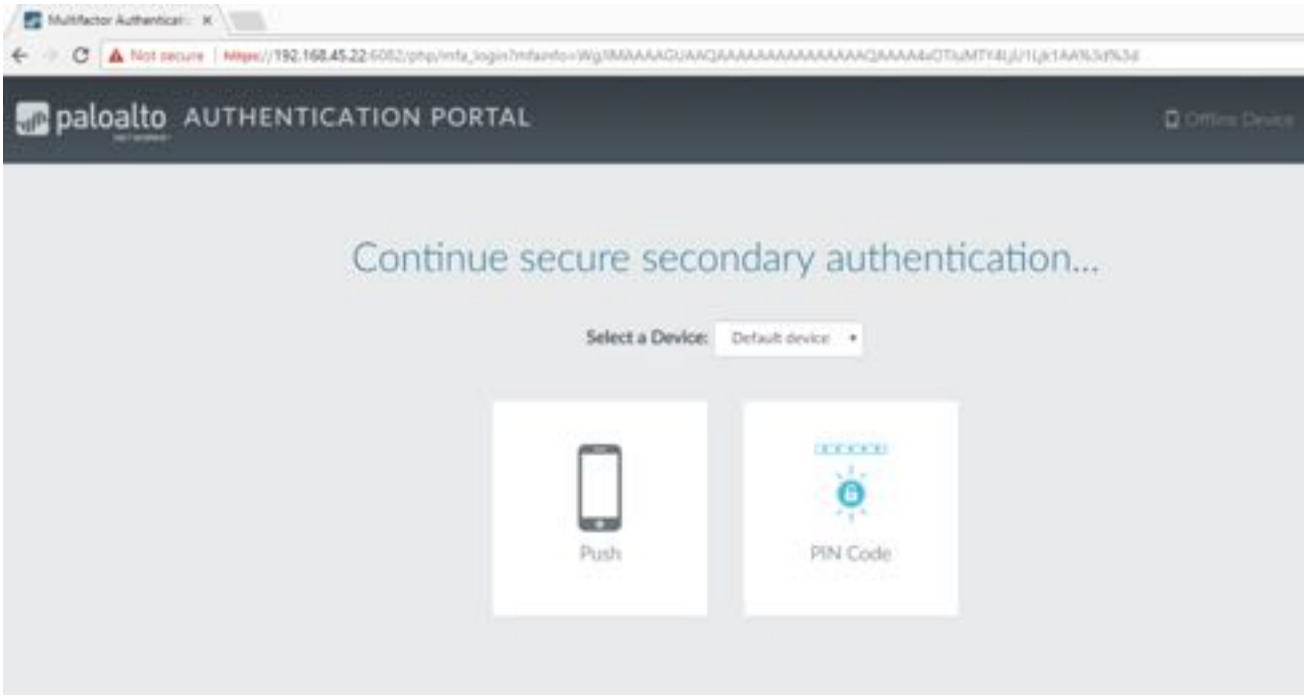


To test an authentication failure, Decline the sign-in request on the mobile device.

2. **PIN Code authentication**

© 2018 Palo Alto Networks, Inc.

1. Ask a user on your network to launch a web browser and access a website. The Captive Portal page with the IP address or hostname for the redirect host you defined earlier should display.
2. Verify that the user enters the credentials for the first authentication factor and then continues to the secondary authentication factor, and selects **PIN Code**.



3. Check that a **PIN Code** displays on the RSA SecurID Access application on the user's mobile device.



4. Ask the user to copy the PIN code in the **Enter the PIN...** prompt of the web browser and click **Submit**. Wait for a few seconds for the firewall to receive the notification of successful authentication. The user should be able to access the requested website.

Configure MFA Between Okta and the Firewall

Multi-factor authentication allows you to protect company assets by using multiple factors to verify the identity of users before allowing them to access network resources.

© 2018 Palo Alto Networks, Inc.

To enable multi-factor authentication (MFA) between the firewall and the Okta identity management service:

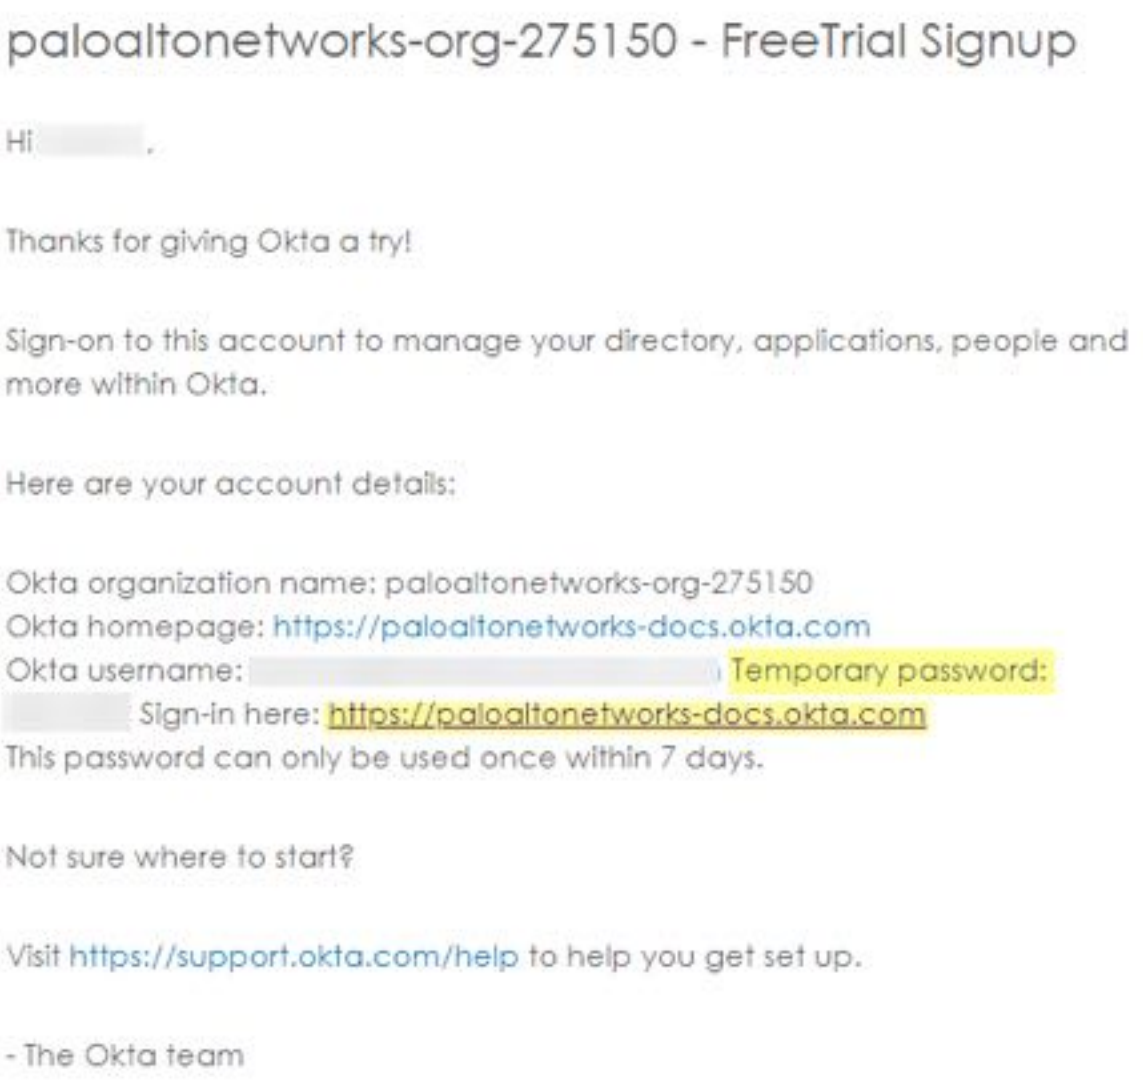
- Configure Okta
- Configure the firewall to integrate with Okta
- Verify MFA with Okta

Configure Okta

Log in to the Okta Admin Portal to create your user accounts, define your Okta MFA policy, and obtain the token information required to configure MFA with Okta on the firewall.

STEP 1 | Create your Okta Admin user account.

1. Submit your email address and name, then click **Get Started**.
2. Click the link in the confirmation email and use the included temporary password to log in to the Okta Admin Portal.


paloaltonetworks-org-275150 - FreeTrial Signup

Hi ,

Thanks for giving Okta a try!

Sign-on to this account to manage your directory, applications, people and more within Okta.

Here are your account details:

Okta organization name: paloaltonetworks-org-275150
Okta homepage: https://paloaltonetworks-docs.okta.com
Okta username: Temporary password:
Sign-in here: https://paloaltonetworks-docs.okta.com
This password can only be used once within 7 days.

Not sure where to start?

Visit https://support.okta.com/help to help you get set up.

- The Okta team

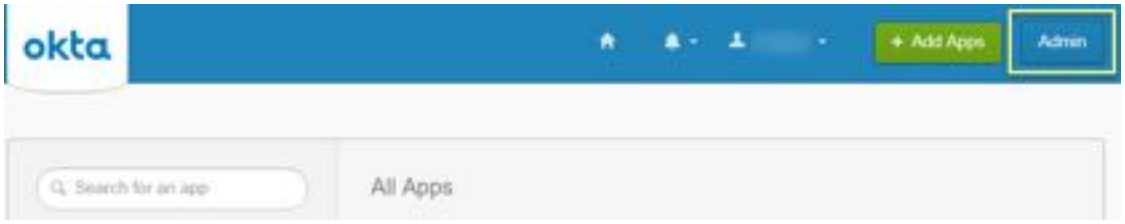
3. Create a new password that includes at least 8 characters, one lowercase letter, one uppercase letter, a number, and does not include any part of your username.
4. Select a password reminder question and enter the answer.

© 2018 Palo Alto Networks, Inc.

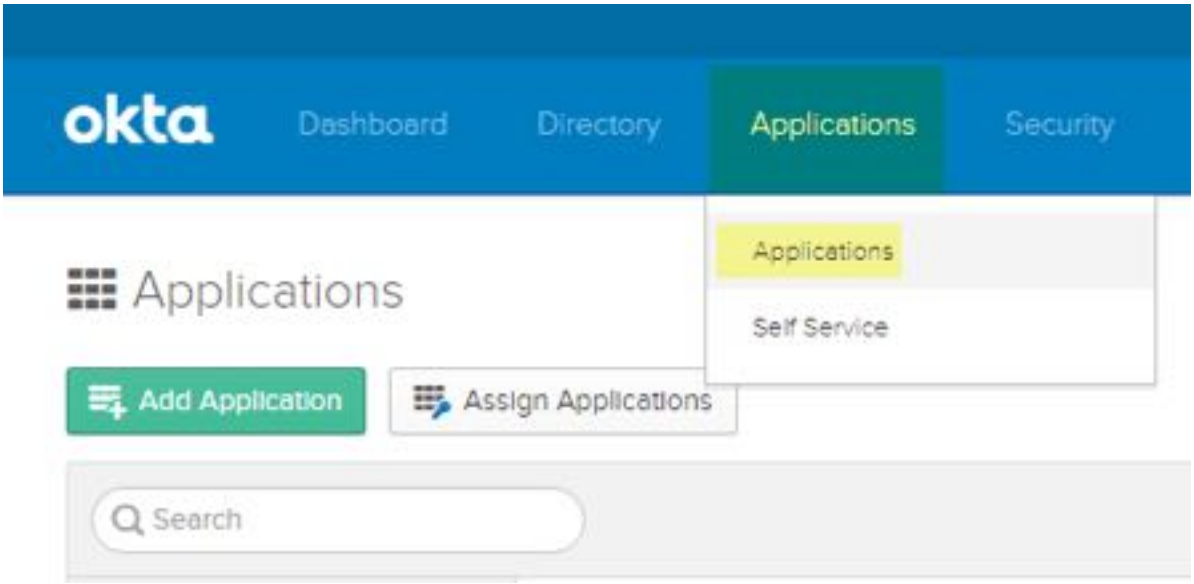
5. Select a security image, then **Create My Account**.

STEP 2 | Configure your Okta service.

If you log in and are not redirected to the Okta Admin Portal, select Admin at the upper right.



1. From the Okta Dashboard, log in with your Okta Admin credentials, then select **Applications** > **Applications**.



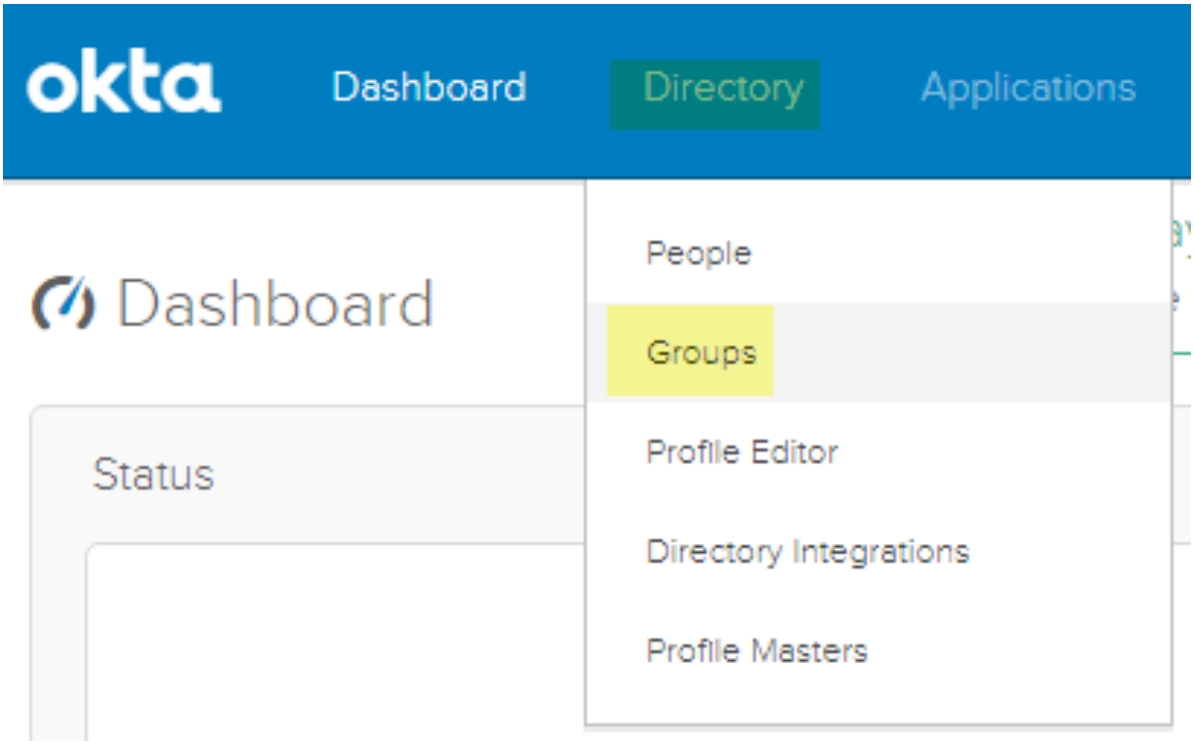
2. Select **Add Application**.
3. Search for **`Okta Verify`**.
4. Select **Add**, then **Done**.



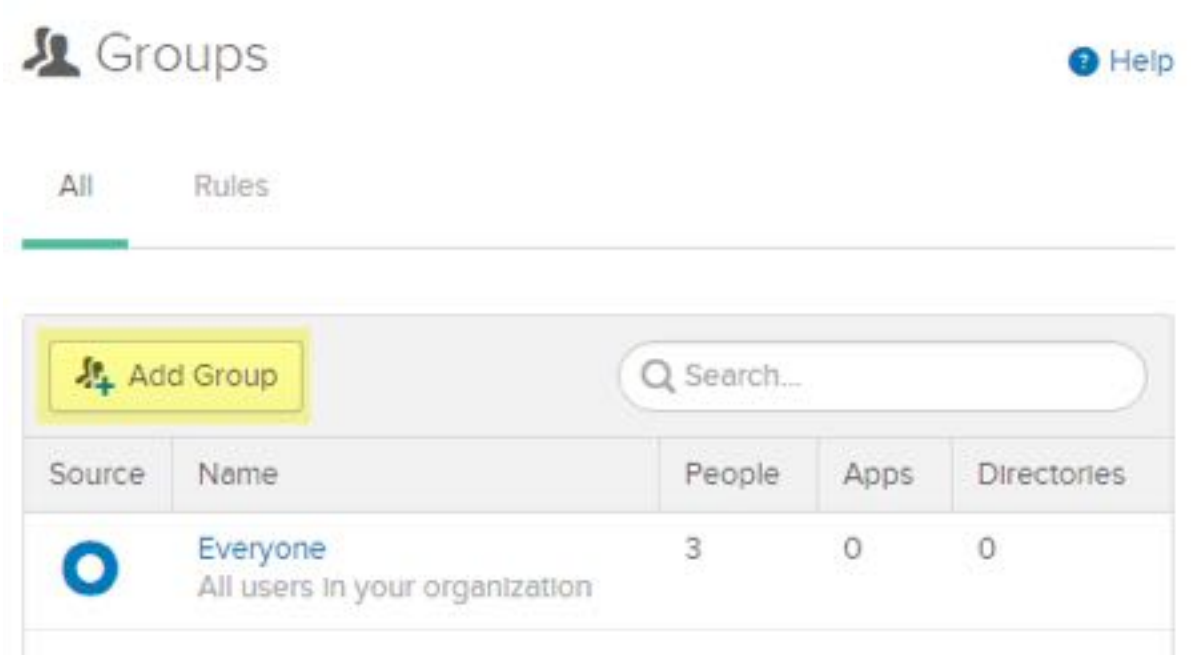
STEP 3 | Create one or more user groups to categorize your users (for example, by device, by policy, or by department) and assign the Okta Verify application.

© 2018 Palo Alto Networks, Inc.

1. Select **Directory** > **Groups**.

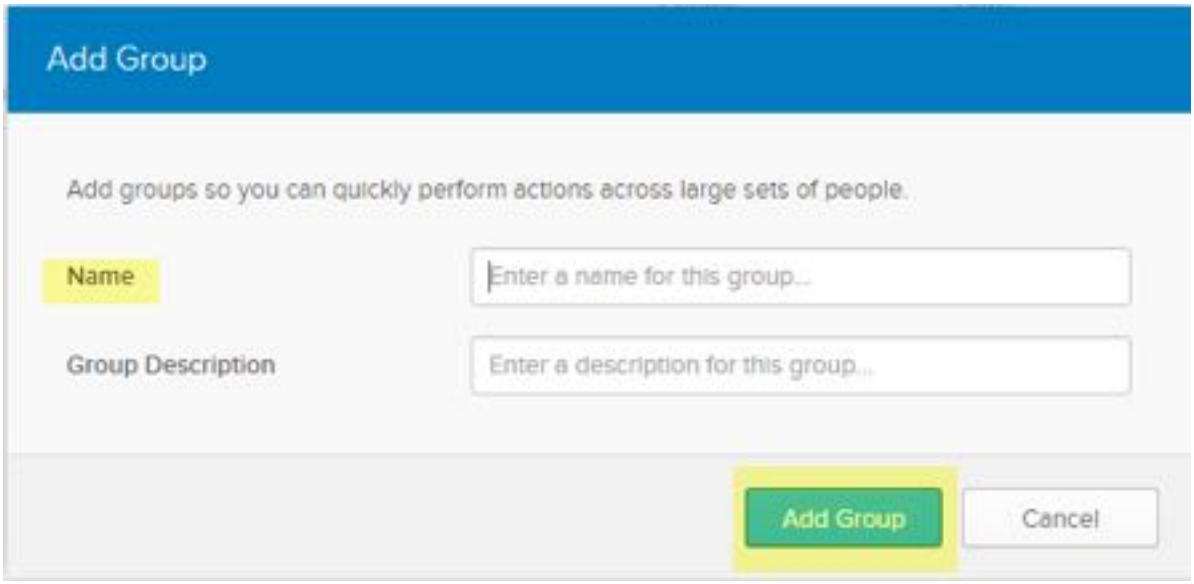


2. Click **Add Group**.



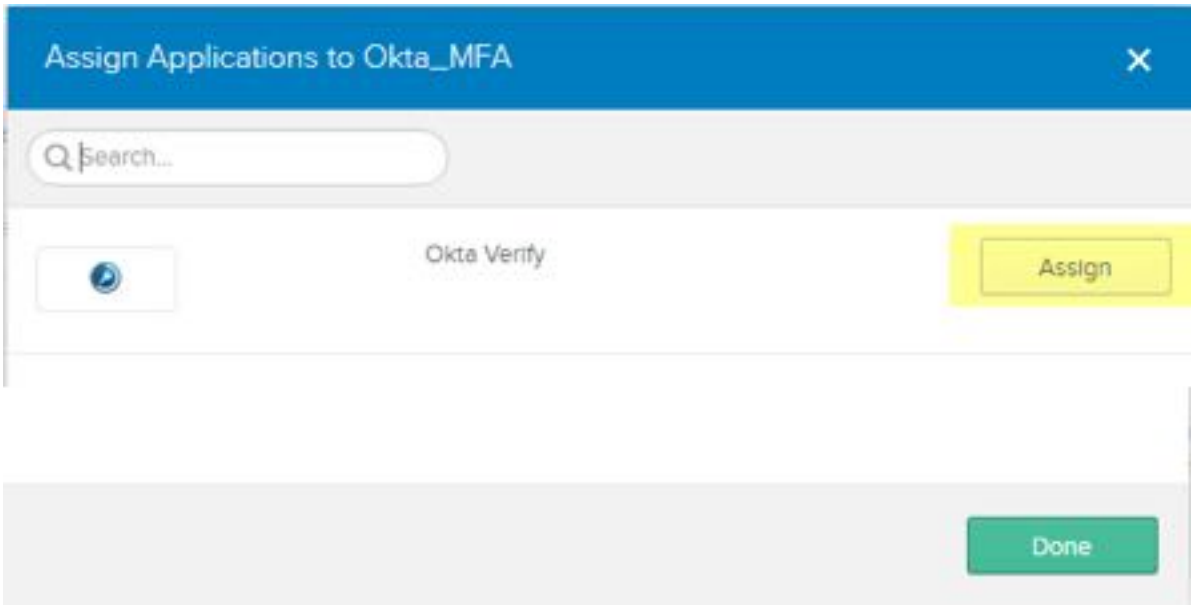
3. Enter a group **Name** and optionally a **Group Description**, then **Add Group**.

© 2018 Palo Alto Networks, Inc.



*The default group **Everyone** includes all users configured for your organization during the first step in* Configure Okta.

4. Select the group you created, then select **Manage Apps**.
5. **Assign** the Okta Verify application you added in Step 2.

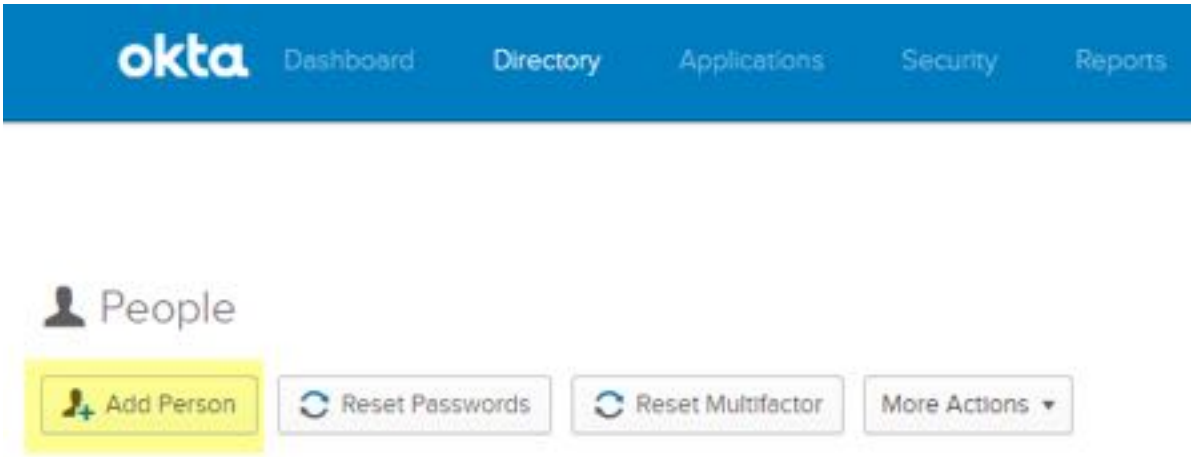


6. After the application has been **Assigned**, click **Done**.
7. Repeat this process for all groups that will use the Okta Verify application for MFA.

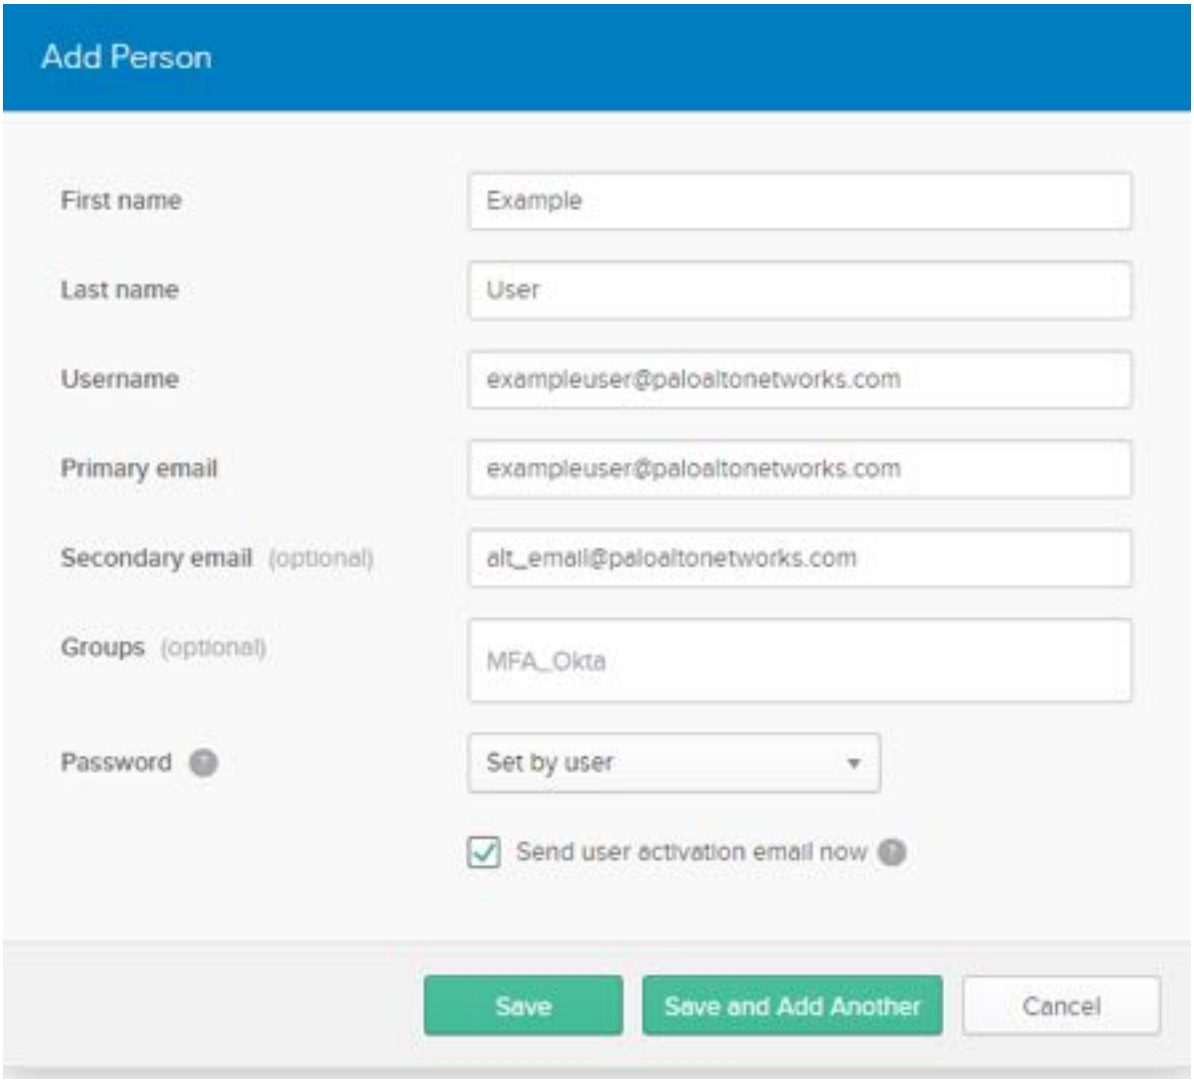
STEP 4 | Add users and assign them to a group.

1. From the Okta Dashboard, select **Directory** > **People** > **Add Person**.

© 2018 Palo Alto Networks, Inc.



2. Enter the user's **First Name**, **Last Name**, and **Username**. The username must match the **Primary email**, which populates automatically, and the username entered on the firewall. You can optionally enter an alternate email address for the user as the **Secondary Email**.



3. Enter the name of the group or **Groups** to associate with this user. When you start typing, the group name populates automatically.

© 2018 Palo Alto Networks, Inc.

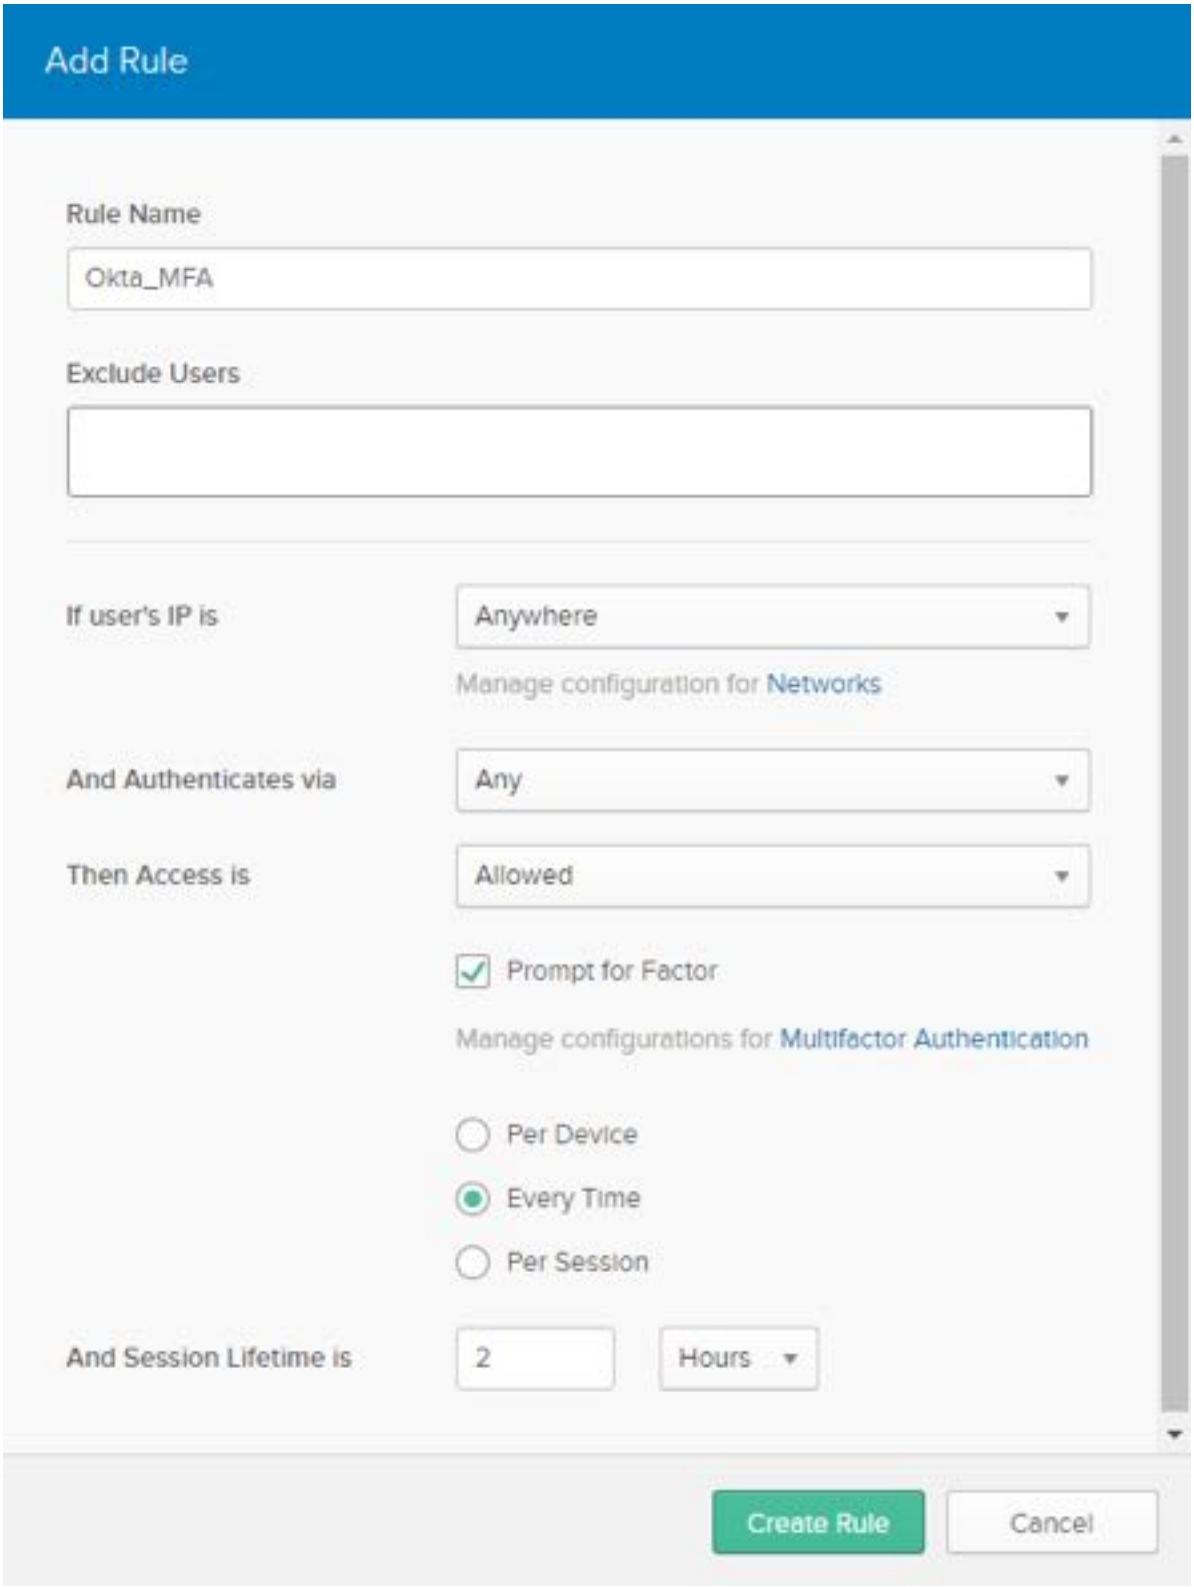
4. Check **Send user activation email now**, then **Save** to add a single user or **Save and Add Another** to continue adding users.

STEP 5 | Assign a test policy to users.

1. Select **Security** > **Authentication** > **Sign On**.

   There is a **Default Policy** with a **Default Rule** that does not prompt users to log in with MFA.
2. Enter the **Rule Name** and check **Prompt for Factor** to enforce the MFA prompt, and select the type of prompt (**Per Device**, **Every Time**, or **Per Session**), then **Create Rule**.

© 2018 Palo Alto Networks, Inc.

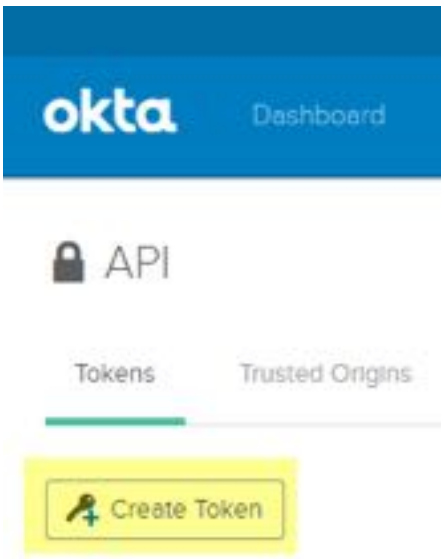


STEP 6 | Record the Okta authentication token information in a safe place because it is only displayed once.

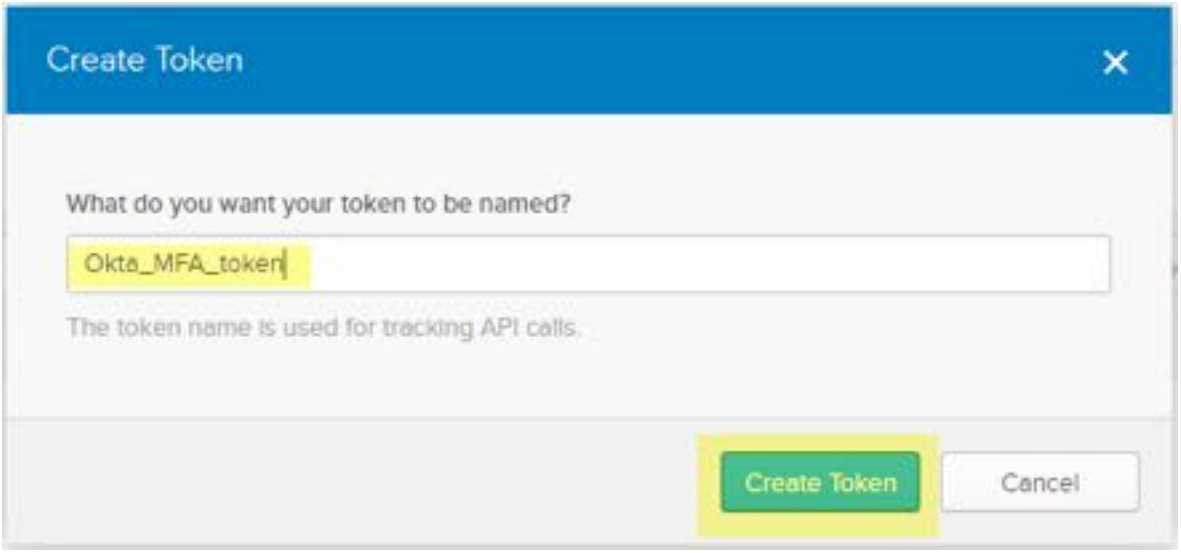
1. Select **Security** > **API** > **Tokens**.

© 2018 Palo Alto Networks, Inc.

2. Select **Create Token**.

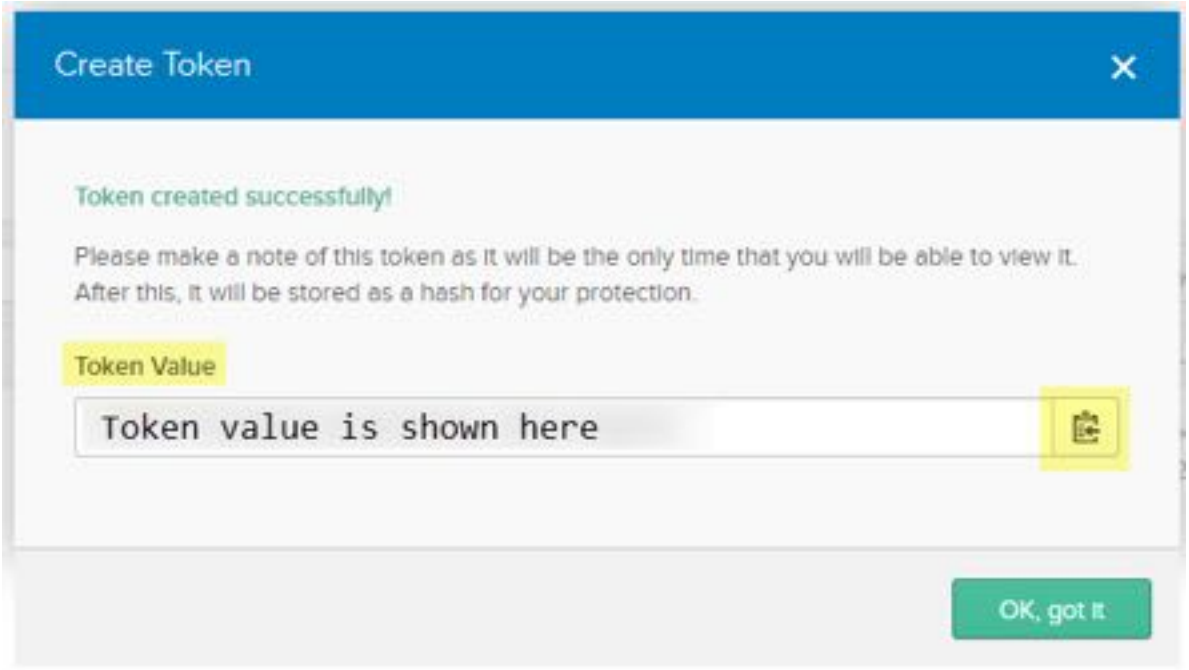


3. Enter a name for the token, then **Create Token**.

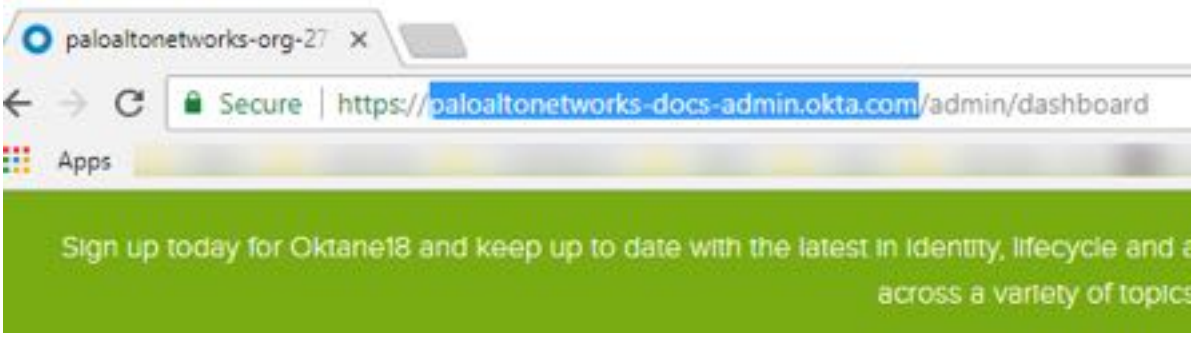


4. Copy the **Token Value**.
   You can click the **Copy to clipboard** button to copy the Token Value to your clipboard.

© 2018 Palo Alto Networks, Inc.



5. In the URL for the Okta Admin Dashboard, copy the portion of the URL after **`https://`** up to **`/admin`** to use as the **API host**.



6. Omit the domain **`okta.com`** from this URL to use as the **Organization**.



For example, in the example Okta Admin Dashboard URL above, **`https://paloaltonetworks-doc-admin.okta.com/admin/dashboard`**:

- The API hostname is **`paloaltonetworks-doc-admin.okta.com`**.
- The Organization is **`paloaltonetworks-doc-admin`**.

STEP 7 | Export the Okta certificates using Base-64 encoding:

1. Depending on your browser, use one of the following methods to export the Okta certificates.
   - **Chrome**—Press **`F12`**, then select **Security** > **View Certificate** > **Details** > **Copy to File**.
   - **Firefox**—Select **Options** > **Privacy & Security** > **View Certificates** > **Export**.
   - **Internet Explorer**—Select **Settings** > **Internet Options** > **Content** > **Certificates** > **Export**.

© 2018 Palo Alto Networks, Inc.

2. Use the Certificate Export Wizard to export the certificates and select **Base-64 encoded X.509** as the format.

Configure the firewall to integrate with Okta

As a prerequisite, confirm that you have mapped all users that you want to authenticate using Okta.

STEP 1 | Import the certificates on the firewall and add them to a Certificate Profile.

STEP 2 | Add a **Multi Factor Authentication Server Profile** for Okta.

1. Select **Device** > **Server Profiles** > **Multi Factor Authentication**.
2. **Add** an MFA server profile.

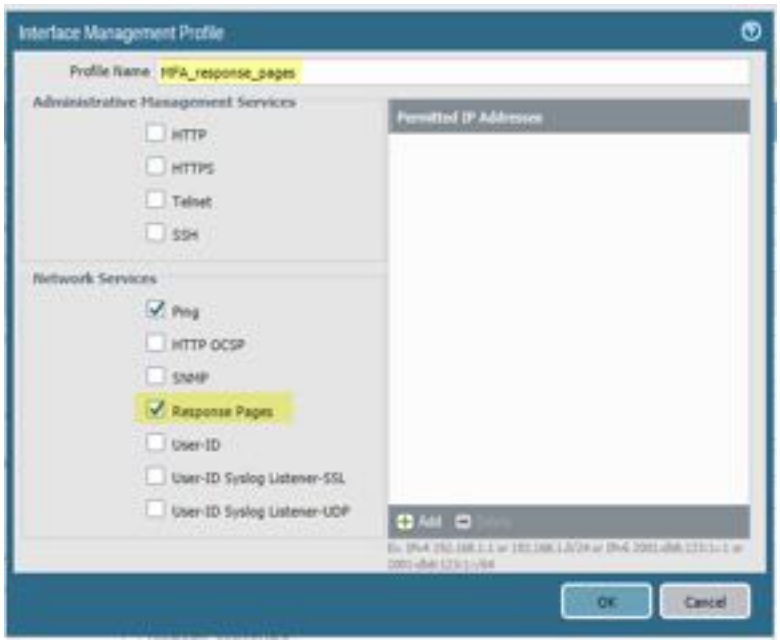


3. Enter a **Profile Name**.
4. Select the **Certificate Profile** you created in Step 1 in Configure the firewall to integrate with Okta.
5. Select **Okta Adaptive** as the **MFA Vendor**.
6. Enter the **API Host**, **Token**, and **Organization** from Step 4 in Configure the firewall to integrate with Okta.

STEP 3 | Configure Captive Portal using **Redirect Mode** to redirect users to the MFA vendor's challenge.

STEP 4 | Enable response pages on the Interface Management Profile to redirect users to the response page challenge.

© 2018 Palo Alto Networks, Inc.



STEP 5 | Create an Authentication Profile and add the MFA vendor as a **Factor** (see Configure Multi-Factor Authentication, Step 3.)



STEP 6 | Enable User-ID on the source zone to require identified users to respond to the challenge using your MFA vendor.

STEP 7 | Create an Authentication Enforcement Object to use the MFA vendor and create an Authentication policy rule (see Configure Authentication Policy, Steps 4 and 5).

STEP 8 | **Commit** your changes.

Verify MFA with Okta

STEP 1 | Verify your users received their enrollment emails, have activated their accounts, and have downloaded the Okta Verify app on their devices.

© 2018 Palo Alto Networks, Inc.

STEP 2 | Go to a website that will prompt the response page challenge.



If you are using a self-signed certificate instead of a PKI-assigned certificate from your organization, a security warning displays that users must click through to access the challenge.

STEP 3 | Log in to the response page using your Okta credentials.

STEP 4 | Confirm the device receives the challenge push notification.

STEP 5 | Confirm users can successfully access the page after authenticating the challenge by accepting the push notification on their devices.

Configure MFA Between Duo and the Firewall

Multi-factor authentication (MFA) allows you to protect company assets by using multiple factors to verify the identity of users before allowing them to access network resources. There are multiple ways to use the Duo identity management service to authenticate with the firewall:

- Two-factor authentication for VPN logins using the GlobalProtect Gateway and a RADIUS server profile (supported on PAN-OS 7.0 and later).
- API-based integration using Captive Portal and an MFA server profile (does not require a Duo Authentication Proxy or SAML IdP - supported on PAN-OS 8.0 and later).
- SAML integration for on-premise servers (supported on PAN-OS 8.0 and later).

To enable SAML MFA between the firewall and Duo to secure administrative access to the firewall:

- Configure Duo for SAML MFA with Duo Access Gateway
- Configure the Firewall to Integrate with Duo
- Verify MFA with Duo

Configure Duo for SAML MFA with Duo Access Gateway

Before you begin, verify that you have deployed the DuoAccessGateway (DAG) on an on-premise server in your DMZ zone.

Create your Duo administrator account and configure the Duo Access Gateway to authenticate your users before they can access resources.

STEP 1 | Create your Duo administrator account.

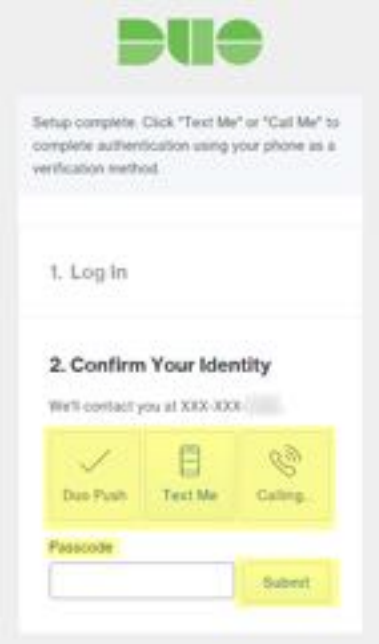
1. On the Duo account creation page, enter your **First Name**, **Last Name**, **Email Address**, **Cell Phone Number**, **Company / Account Name**, and select the number of employees in the organization.
2. Agree to the Terms and Privacy Policy and respond to the reCAPTCHA challenge to **Create My Account**.

© 2018 Palo Alto Networks, Inc.



STEP 2 | Verify your Duo administrator account.

1. Select the authentication verification method (**Duo Push**, **Text Me**, or **Calling...**).
2. Enter the **Passcode** you receive and **Submit** it to verify your account.

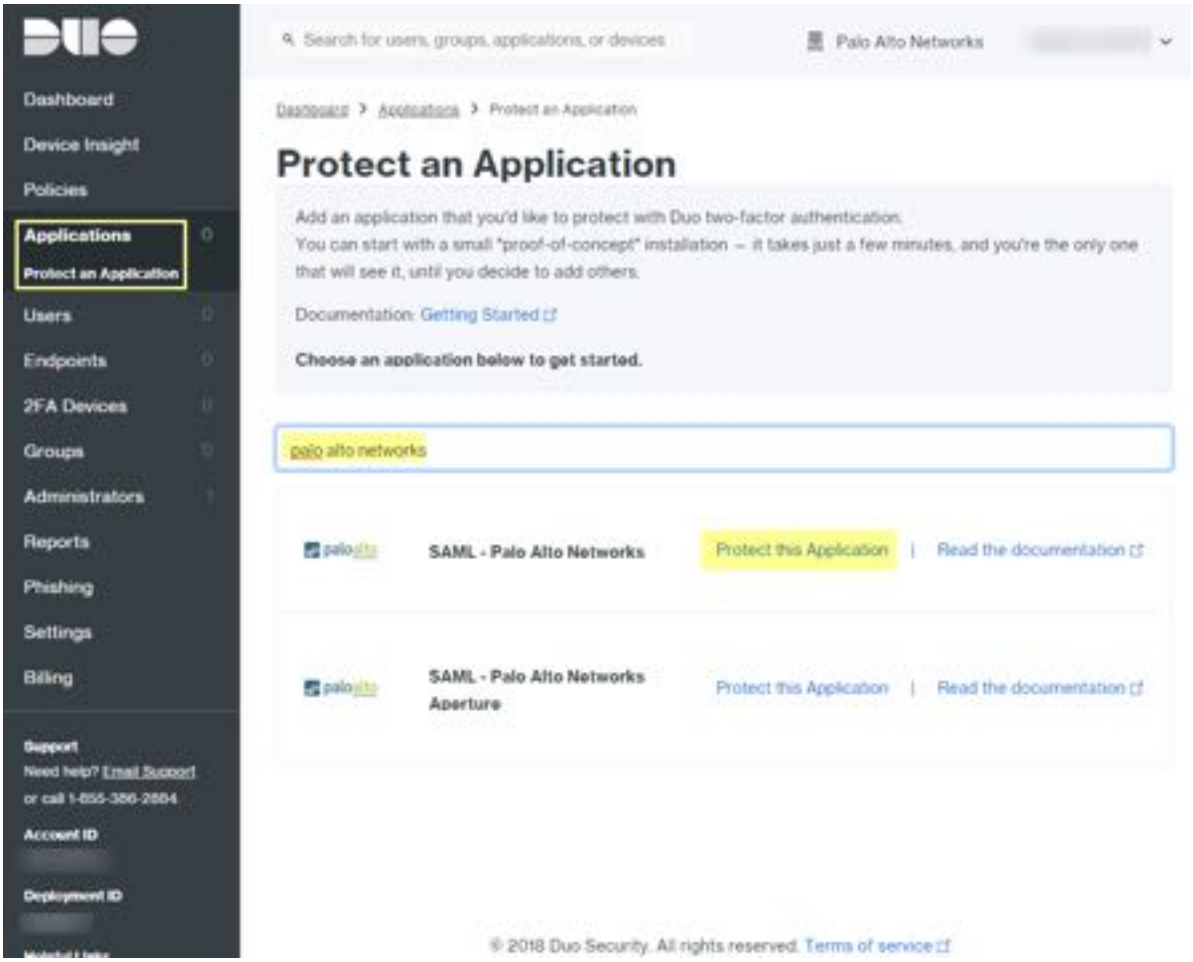


STEP 3 | Configure your Duo service for SAML.

After creating your configuration, download the configuration file at the top of the page.

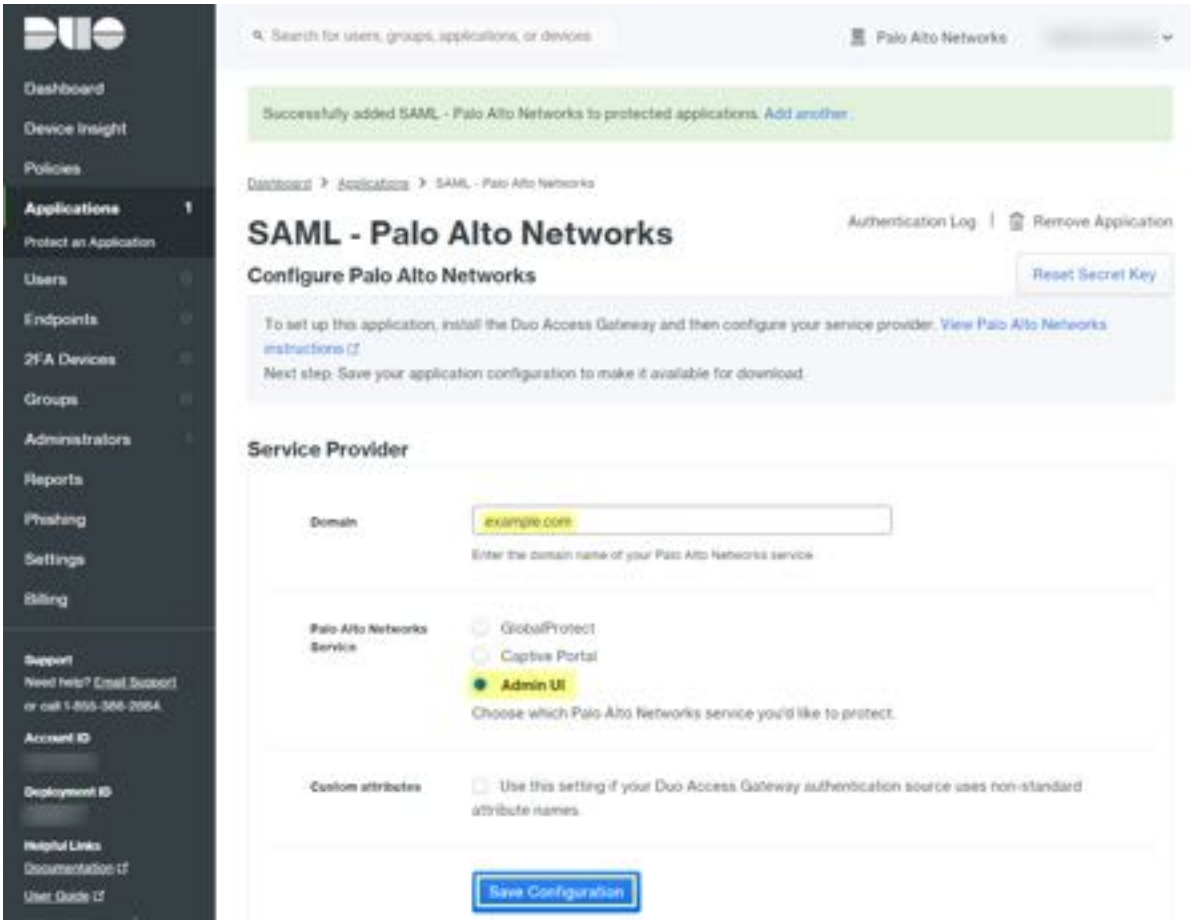
1. In the Duo Admin Panel, select **Applications** > **Protect an Application**.
2. Enter **`Palo Alto Networks`** to search the applications.
3. Locate **SAML - Palo Alto Networks** in the list of results, then **Protect this Application**.

© 2018 Palo Alto Networks, Inc.



4. Enter the **Domain**.
5. Select **Admin UI** as the **Palo Alto Networks Service**.
6. Configure your **Policy** and other **Settings**, and **Save Configuration**.

© 2018 Palo Alto Networks, Inc.



7. **Download your configuration file**.

   The link to download the file is at the top of the page.



STEP 4 | Upload the configuration file to the Duo Access Gateway (DAG).

1. In the DAG admin console, select **Applications**.
2. Click **Choose File** and select the configuration file you downloaded, then **Upload** it.
3. In **Settings** > **Session Management**, disable **User agent binding**, then **Save Settings**.

STEP 5 | In the DAG admin console, configure your Active Directory or OpenLDAP server as the authentication source and download the metadata file.

1. Log in to the DAG admin console.
2. In **Authentication Source** > **Set Active Source**, select your **Source type** (Active Directory or OpenLDAP) and **Set Active Source**.

© 2018 Palo Alto Networks, Inc.

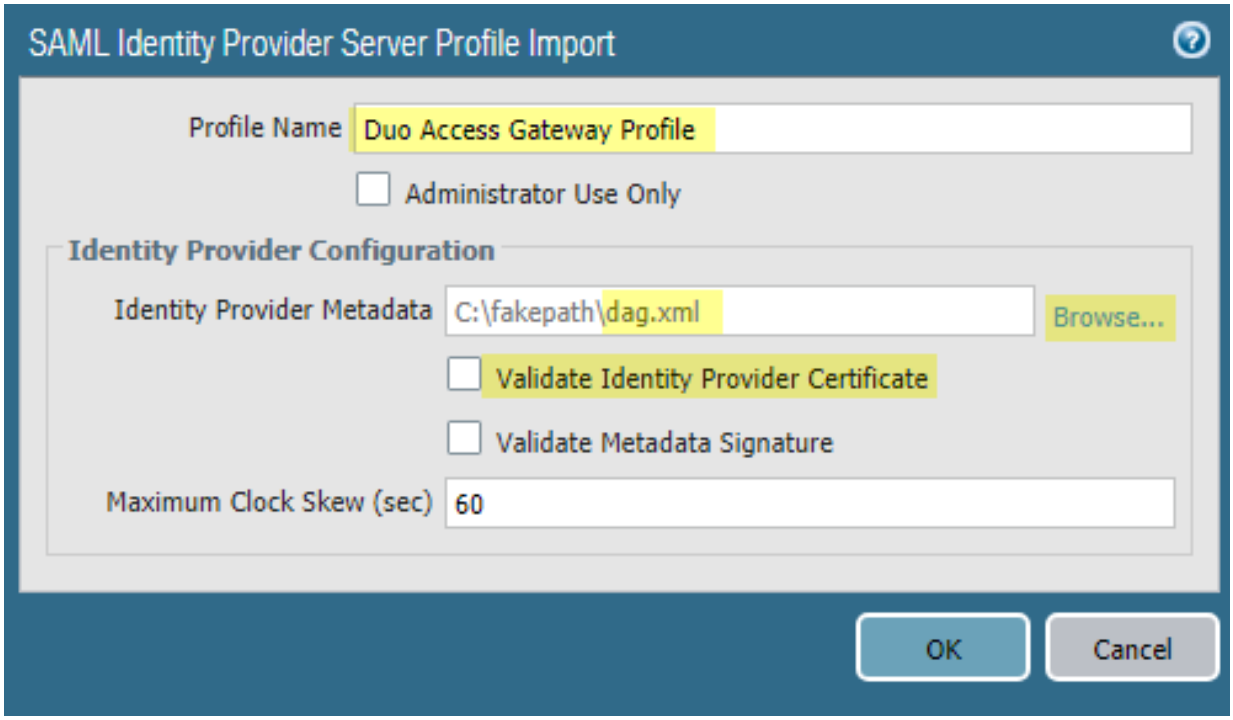
3. In **Configure Sources**, enter the **Attributes**.
   - For Active Directory, enter **`mail,sAMAccountName,userPrincipalName,objectGUID`**.
   - For OpenLDAP, enter **`mail,uid`**.
   - For any custom attributes, append them to the end of the list and separate each attribute with a comma. Do not delete any existing attributes.
4. **Save Settings** to save the configuration.
5. Select **Applications** > **Metadata**, then click **Download XML metadata** to download the XML metadata you will need to import into the firewall.

   The file will be named dag.xml. Because this file includes sensitive information to authenticate your Duo account with the firewall, make sure to keep the file in a secure location to avoid the risk of compromising this information.

Configure the Firewall to Integrate with Duo

STEP 1 | Import the Duo metadata.

1. Log on to the firewall web interface.
2. On the firewall, select **Device** > **Server Profiles** > **SAML Identity Provider** > **Import**.
3. Enter the **Profile Name**.
4. **Browse** to the **Identity Provider Metadata** file (**dag.xml**).
5. Disable **Validate Identity Provider Certificate**, then click **OK**.



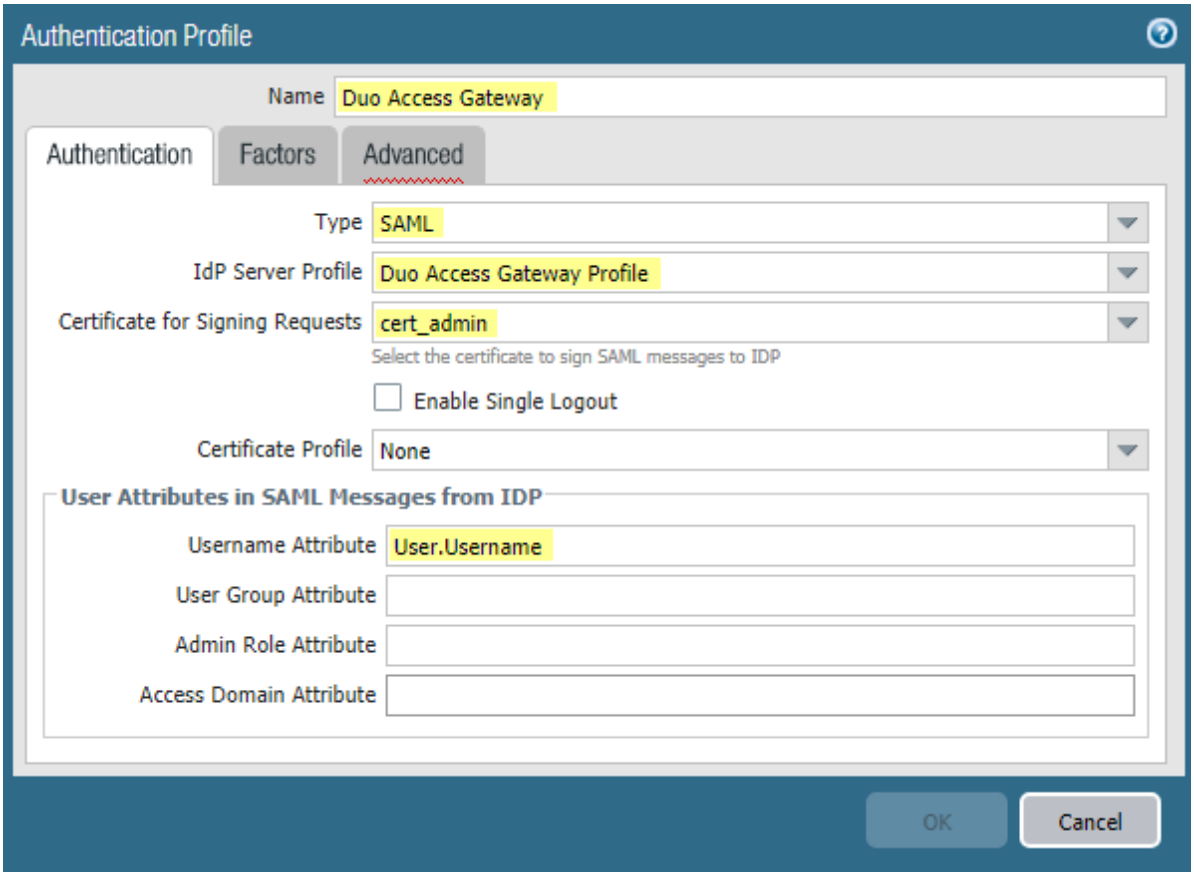
STEP 2 | Add an authentication profile.

The authentication profile allows Duo as the identity provider that validates administrator login credentials.

1. **Add** an **Authentication Profile**.
2. Enter the profile **Name**.
3. Select **SAML** as the authentication **Type**.
4. Select **Duo Access Gateway Profile** as the **IdP Server Profile**.

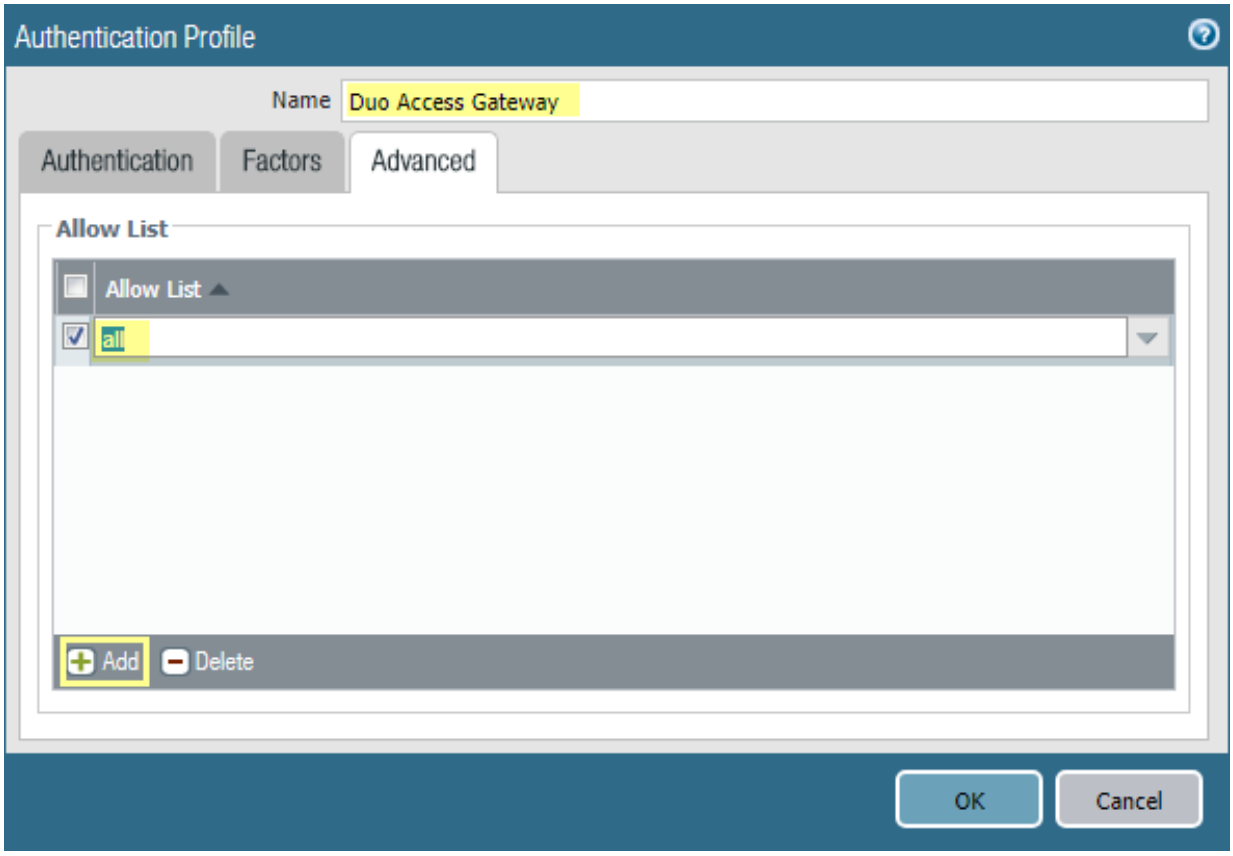
© 2018 Palo Alto Networks, Inc.

5. Select the certificate you want to use for SAML communication with the Duo Access Gateway for the **Certificate for Signing Requests**.
6. Enter **`user.username`** as the **Username Attribute**.



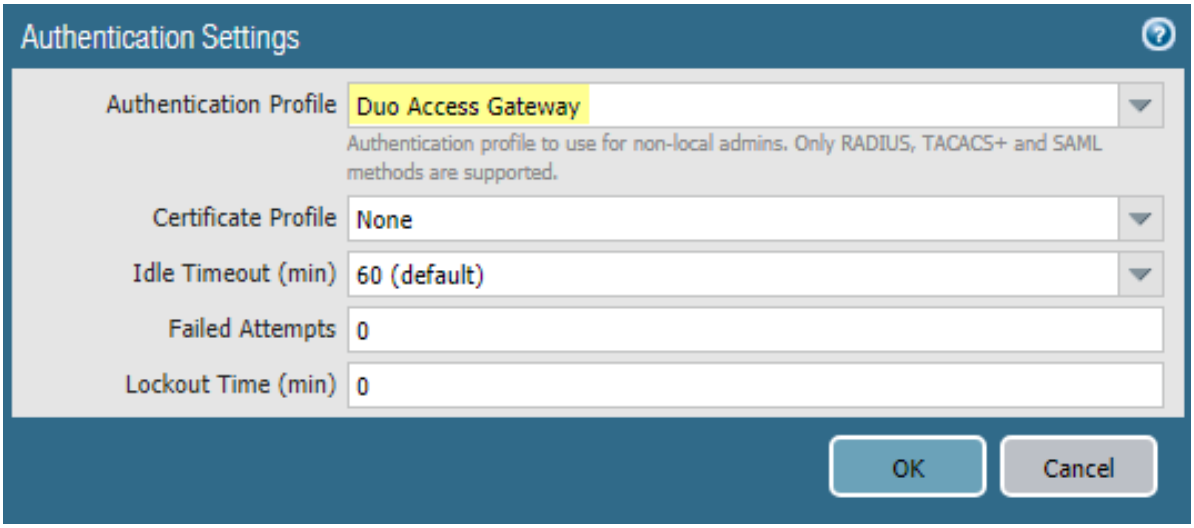
7. Select **Advanced** to **Add** an allow list.
8. Select **all**, then click **OK**.
9. **Commit** the changes.

© 2018 Palo Alto Networks, Inc.



STEP 3 | Specify the authentication settings that the firewall uses for SAML authentication with Duo.

1. Select **Device** > **Setup** > **Management** and edit the **Authentication Settings**.
2. Select **Duo Access Gateway** as the **Authentication Profile**, then click **OK**.



3. **Commit** your changes.

STEP 4 | Add accounts for administrators who will authenticate to the firewall using Duo.

1. Select **Device** > **Administrators** and **Add** an account.
2. Enter a user **Name**.
3. Select **Duo Access Gateway** as the **Authentication Profile**.

© 2018 Palo Alto Networks, Inc.

4. Select the **Administrator Type**, then click **OK**.

   Select **Role Based** if you want to use a custom role for the user. Otherwise, select **Dynamic**. To require administrators to log in using SSO with Duo, assign the authentication profile to all current administrators.



Verify MFA with Duo

STEP 1 | Log in to the web interface on the firewall.

STEP 2 | Select **Use Single Sign-On** and **Continue**.

STEP 3 | Enter your login credentials on the Duo Access Gateway login page.

STEP 4 | Select an authentication method (push notification, phone call, or passcode entry).

When you authenticate successfully, you will be redirected to the firewall web interface.

© 2018 Palo Alto Networks, Inc.

Configure SAML Authentication

To configure SAML single sign-on (SSO) and single logout (SLO), you must register the firewall and the IdP with each other to enable communication between them. If the IdP provides a metadata file containing registration information, you can import it onto the firewall to register the IdP and to create an IdP server profile. The server profile defines how to connect to the IdP and specifies the certificate that the IdP uses to sign SAML messages. You can also use a certificate for the firewall to sign SAML messages. Using certificates is optional but recommended to secure communications between the firewall and the IdP.

Palo Alto Networks requires HTTPS to ensure the confidentiality of all SAML transactions instead of alternative approaches such as encrypted SAML assertions. To ensure the integrity of all messages processed in a SAML transaction, Palo Alto Networks strongly recommends requiring digital certificates to cryptographically sign all messages.

The following procedure describes how to configure SAML authentication for end users and firewall administrators. You can also configure SAML authentication for Panorama administrators.



SSO is available to administrators and to GlobalProtect and Captive Portal end users. SLO is available to administrators and GlobalProtect end users, but not to Captive Portal end users.

Administrators can use SAML to authenticate to the firewall web interface, but not to the CLI.

STEP 1 | (Recommended) Obtain the certificates that the IdP and firewall will use to sign SAML messages.

If the certificates don't specify key usage attributes, all usages are allowed by default, including signing messages. In this case, you can Obtain Certificates by any method.

If the certificates do specify key usage attributes, one of the attributes must be Digital Signature, which is not available on certificates that you generate on the firewall or Panorama. In this case, you must import the certificates:

- **Certificate the firewall uses to sign SAML messages**—Import the certificate from your enterprise certificate authority (CA) or a third-party CA.
- **Certificate the IdP uses to sign SAML messages**—Import a metadata file containing the certificate from the IdP (see the next step). The IdP certificate is limited to the following algorithms:

  Public key algorithms—RSA (1,024 bits or larger) and ECDSA (all sizes). A firewall in FIPS/CC mode supports RSA (2,048 bits or larger) and ECDSA (all sizes).

  Signature algorithms—SHA1, SHA256, SHA384, and SHA512. A firewall in FIPS/CC mode supports SHA256, SHA384, and SHA512.

STEP 2 | Add a SAML IdP server profile.

The server profile registers the IdP with the firewall and defines how they connect.

In this example, you import a SAML metadata file from the IdP so that the firewall can automatically create a server profile and populate the connection, registration, and IdP certificate information.



If the IdP doesn't provide a metadata file, select Device > Server Profiles > SAML Identity Provider, Add the server profile, and manually enter the information (consult your IdP administrator for the values).

1. Export the SAML metadata file from the IdP to a client system that the firewall can access.

   The certificate specified in the file must meet the requirements listed in the preceding step. Refer to your IdP documentation for instructions on exporting the file.

© 2018 Palo Alto Networks, Inc.

2. Select **Device** > **Server Profiles** > **SAML Identity Provider** and **Import** the metadata file onto the firewall.
3. Enter a **Profile Name** to identify the server profile.
4. **Browse** to the **Identity Provider Metadata** file.
5. (Recommended) Select **Validate Identity Provider Certificate** (default) to have the firewall validate the **Identity Provider Certificate**.

   Validation occurs only after you assign the server profile to an authentication profile and **Commit**. The firewall uses the **Certificate Profile** in the authentication profile to validate the certificate.

   Validating the certificate is a best practice for improved security.
6. Enter the **Maximum Clock Skew**, which is the allowed difference in seconds between the system times of the IdP and the firewall at the moment when the firewall validates IdP messages (default is 60; range is 1 to 900). If the difference exceeds this value, authentication fails.
7. Click **OK** to save the server profile.
8. Click the server profile Name to display the profile settings. Verify that the imported information is correct and edit it if necessary.

STEP 3 | Configure an authentication profile.

The profile defines authentication settings that are common to a set of users.

1. Select **Device** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Set the **Type** to **SAML**.
4. Select the **IdP Server Profile** you configured.
5. Select the **Certificate for Signing Requests**.

   The firewall uses this certificate to sign messages it sends to the IdP. You can import a certificate generated by your enterprise CA or you can generate a certificate using the root CA that was generated on the firewall or Panorama.
6. (**Optional**) **Enable Single Logout** (disabled by default).
7. Select the **Certificate Profile** that the firewall will use to validate the **Identity Provider Certificate**.
8. Enter the **Username Attribute** that IdP messages use to identify users (default **`username`**).

   If you manage administrator authorization in the IdP identity store, specify the Admin Role Attribute and Access Domain Attribute also.
9. Select **Advanced** and **Add** the users and user groups that are allowed to authenticate with this authentication profile.
10. Click **OK** to save the authentication profile.

STEP 4 | Assign the authentication profile to firewall applications that require authentication.

1. Assign the authentication profile to:
   - Administrator accounts that you manage locally on the firewall. In this example, Configure a Firewall Administrator Account before you verify the SAML configuration later in this procedure.
   - Administrator accounts that you manage externally in the IdP identity store. Select **Device** > **Setup** > **Management**, edit the Authentication Settings, and select the **Authentication Profile** you configured.
   - Authentication policy rules that secure the services and applications that end users access through Captive Portal. See Configure Authentication Policy.
   - GlobalProtect portals and gateways that end users access.
2. **Commit** your changes.

© 2018 Palo Alto Networks, Inc.

The firewall validates the **Identity Provider Certificate** that you assigned to the SAML IdP server profile.

STEP 5 | Create a SAML metadata file to register the firewall application (management access, Captive Portal, or GlobalProtect) on the IdP.

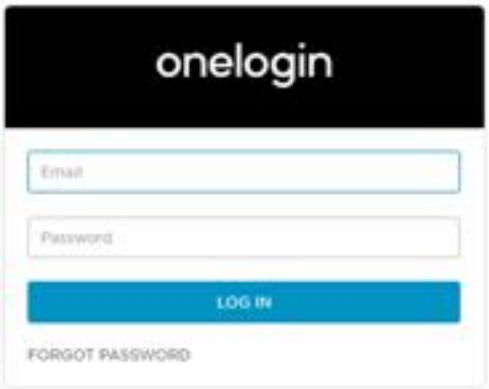
1. Select **Device** > **Authentication Profile** and, in the Authentication column for the authentication profile you configured, click **Metadata**.
2. In the **Commands** drop-down, select the application you want to register:
   - **management** (default)—Administrative access to the web interface.
   - **captive-portal**—End user access to services and applications through Captive Portal.
   - **global-protect**—End user access to services and applications through GlobalProtect.
3. (Captive Portal or GlobalProtect only) for the **Vsysname Combo**, select the virtual system in which the Captive Portal settings or GlobalProtect portal are defined.
4. Enter the interface, IP address, or hostname based on the application you will register:
   - **management**—For the **Management Choice**, select **Interface** (default) and select an interface that is enabled for management access to the web interface. The default selection is the IP address of the MGT interface.
   - **captive-portal**—For the **IP Hostname**, enter the IP address or hostname of the **Redirect Host** (see **Device** > **User Identification** > **Captive Portal Settings**).
   - **global-protect**—For the **IP Hostname**, enter the hostname or IP address of the GlobalProtect portal or gateway.
5. Click **OK** and save the metadata file to your client system.
6. Import the metadata file into the IdP server to register the firewall application. Refer to your IdP documentation for instructions.

STEP 6 | Verify that users can authenticate using SAML SSO.

For example, to verify that SAML is working for access to the web interface using a local administrator account:

1. Go to the URL of the firewall web interface.
2. Click **Use Single Sign-On**.
3. Enter the username of the administrator.
4. Click **Continue**.

   The firewall redirects you to authenticate to the IdP, which displays a login page. For example:



5. Log in using your SSO username and password.

   After you successfully authenticate on the IdP, it redirects you back to the firewall, which displays the web interface.

© 2018 Palo Alto Networks, Inc.

6. Use your firewall administrator account to request access to another SSO application.
   Successful access indicates SAML SSO authentication succeeded.

© 2018 Palo Alto Networks, Inc.

Configure Kerberos Single Sign-On

Palo Alto Networks firewalls and Panorama support Kerberos V5 single sign-on (SSO) to authenticate administrators to the web interface and end users to Captive Portal. With Kerberos SSO enabled, the user needs to log in only for initial access to your network (such as logging in to Microsoft Windows). After this initial login, the user can access any browser-based service in the network (such as the firewall web interface) without having to log in again until the SSO session expires.

STEP 1 | Create a Kerberos keytab.

The keytab is a file that contains the principal name and password of the firewall, and is required for the SSO process.



If the Kerberos SSO hostname is included in the request sent to the firewall, then the hostname must match the service principal name of the keytab; otherwise, the Kerberos authentication request is not sent.

1. Create Kerberos account for the firewall. Refer to your Kerberos documentation for the steps.
2. Log in to the KDC and open a command prompt.
3. Enter the following command, where *<principal_name>*, *<password>*, and *<algorithm>* are variables.

```
ktpass /princ <principal_name> /pass <password> /crypto <algorithm> /ptype
 KRB5_NT_PRINCIPAL /out <file_name>.keytab
```



If the firewall is in FIPS/CC mode, the algorithm must be ***aes128-cts-hmac-sha1-96*** *or* ***aes256-cts-hmac-sha1-96****. Otherwise, you can also use* ***des3-cbc-sha1 or arcfour-hmac****. To use an Advanced Encryption Standard (AES) algorithm, the functional level of the KDC must be Windows Server 2008 or later and you must enable AES encryption for the firewall account.*

The algorithm in the keytab must match the algorithm in the service ticket that the TGS issues to clients. Your Kerberos administrator determines which algorithms the service tickets use.

STEP 2 | Configure an Authentication Profile and Sequence to define Kerberos settings and other authentication options that are common to a set of users.

- Enter the **Kerberos Realm** (usually the DNS domain of the users, except that the realm is uppercase).
- **Import** the **Kerberos Keytab** that you created for the firewall.

STEP 3 | Assign the authentication profile to the firewall application that requires authentication.

- Administrative access to the web interface—Configure a Firewall Administrator Account and assign the authentication profile you configured.
- End user access to services and applications—Assign the authentication profile you configured to an authentication enforcement object. When configuring the object, set the **Authentication Method** to **browser-challenge**. Assign the object to Authentication policy rules. For the full procedure to configure authentication for end users, see Configure Authentication Policy.

© 2018 Palo Alto Networks, Inc.

Configure Kerberos Server Authentication

You can use Kerberos to natively authenticate end users and firewall or Panorama administrators to an Active Directory domain controller or a Kerberos V5-compliant authentication server. This authentication method is interactive, requiring users to enter usernames and passwords.

To use a Kerberos server for authentication, the server must be accessible over an IPv4 address. IPv6 addresses are not supported.

STEP 1 | Add a Kerberos server profile.

The profile defines how the firewall connects to the Kerberos server.

1. Select **Device** > **Server Profiles** > **Kerberos** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** each server and specify a **Name** (to identify the server), IPv4 address or FQDN of the **Kerberos Server**, and optional **Port** number for communication with the server (default 88).

   If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change in order for the new server address to take effect.
4. Click **OK** to save your changes to the profile.

STEP 2 | Assign the server profile to an Configure an Authentication Profile and Sequence.

The authentication profile defines authentication settings that are common to a set of users.

STEP 3 | Assign the authentication profile to the firewall application that requires authentication.

- Administrative access to the web interface—Configure a Firewall Administrator Account and assign the authentication profile you configured.
- End user access to services and applications—Assign the authentication profile you configured to an authentication enforcement object and assign the object to Authentication policy rules. For the full procedure to configure authentication for end users, see Configure Authentication Policy.

STEP 4 | Verify that the firewall can Test Authentication Server Connectivity to authenticate users.

© 2018 Palo Alto Networks, Inc.

Configure TACACS+ Authentication

You can configure TACACS+ authentication for end users and firewall or Panorama administrators. You can also use a TACACS+ server to manage administrator authorization (role and access domain assignments) by defining Vendor-Specific Attributes (VSAs). For all users, you must configure a TACACS+ server profile that defines how the firewall or Panorama connects to the server (see Step 1 below). You then assign the server profile to an authentication profile for each set of users who require common authentication settings (see Step 2 below). What you do with the authentication profile depends on which users the TACACS+ server authenticates:

- **End users**—Assign the authentication profile to an authentication enforcement object and assign the object to Authentication policy rules. For the full procedure, see Configure Authentication Policy.
- **Administrative accounts with authorization managed locally on the firewall or Panorama**—Assign the authentication profile to firewall administrator or Panorama administrator accounts.
- **Administrative accounts with authorization managed on the TACACS+ server**—The following procedure describes how to configure TACACS+ authentication and authorization for firewall administrators. For Panorama administrators, refer to Configure TACACS+ Authentication for PanoramaAdministrators.

STEP 1 | Add a TACACS+ server profile.

The profile defines how the firewall connects to the TACACS+ server.

1. Select **Device** > **Server Profiles** > **TACACS+** and **Add** a profile.
2. Enter a **Profile Name** to identify the server profile.
3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
4. Select the **Authentication Protocol** (default is **CHAP**) that the firewall uses to authenticate to the TACACS+ server.

   Select CHAP if the TACACS+ server supports that protocol; it is more secure than PAP.
5. **Add** each TACACS+ server and enter the following:
   - **Name** to identify the server
   - **TACACS+ Server** IP address or FQDN. If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
   - **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords)
   - Server **Port** for authentication requests (default is 49)
6. Click **OK** to save the server profile.

STEP 2 | Assign the TACACS+ server profile to an authentication profile.

The authentication profile defines authentication settings that are common to a set of users.

1. Select **Device** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Set the **Type** to **TACACS+**.
4. Select the **Server Profile** you configured.
5. Select **Retrieve user group from TACACS+** to collect user group information from VSAs defined on the TACACS+ server.

© 2018 Palo Alto Networks, Inc.

The firewall matches the group information against the groups you specify in the Allow List of the authentication profile.

6. Select **Advanced** and, in the Allow List, **Add** the users and groups that are allowed to authenticate with this authentication profile.
7. Click **OK** to save the authentication profile.

STEP 3 | Configure the firewall to use the authentication profile for all administrators.

1. Select **Device** > **Setup** > **Management** and edit the Authentication Settings.
2. Select the **Authentication Profile** you configured and click **OK**.

STEP 4 | Configure the roles and access domains that define authorization settings for administrators.

If you already defined TACACS+ VSAs on the TACACS+ server, the names you specify for roles and access domains on the firewall must match the VSA values.

1. Configure an Admin Role Profile if the administrator will use a custom role instead of a predefined (dynamic) role.
2. Configure an access domain if the firewall has more than one virtual system—Select **Device** > **Access Domain**, **Add** an access domain, enter a **Name** to identify the access domain, and **Add** each virtual system that the administrator will access, and then click **OK**.

STEP 5 | **Commit** your changes to activate them on the firewall.

STEP 6 | Configure the TACACS+ server to authenticate and authorize administrators.

Refer to your TACACS+ server documentation for the specific instructions to perform these steps:

1. Add the firewall IP address or hostname as the TACACS+ client.
2. Add the administrator accounts.

   If you selected CHAP as the Authentication Protocol, you must define accounts with reversibly encrypted passwords. Otherwise, CHAP authentication will fail.
3. Define TACACS+ VSAs for the role, access domain, and user group of each administrator.

STEP 7 | Verify that the TACACS+ server performs authentication and authorization for administrators.

1. Log in the firewall web interface using an administrator account that you added to the TACACS+ server.
2. Verify that you can access only the web interface pages that are allowed for the role you associated with the administrator.
3. In the **Monitor**, **Policies**, and **Objects** tabs, verify that you can access only the virtual systems that are allowed for the access domain you associated with the administrator.

© 2018 Palo Alto Networks, Inc.

Configure RADIUS Authentication

You can configure RADIUS authentication for end users and firewall or Panorama administrators. For administrators, you can use RADIUS to manage authorization (role and access domain assignments) by defining Vendor-Specific Attributes (VSAs). You can also use RADIUS to implement Multi-Factor Authentication (MFA) for administrators and end users. To enable RADIUS authentication, you must configure a RADIUS server profile that defines how the firewall or Panorama connects to the server (see Step 1 below). You then assign the server profile to an authentication profile for each set of users who require common authentication settings (see Step 2 below). What you do with the authentication profile depends on which users the RADIUS server authenticates:

- **End users**—Assign the authentication profile to an authentication enforcement object and assign the object to Authentication policy rules. For the full procedure, see Configure Authentication Policy.



You can also configure client systems to send RADIUS Vendor-SpecificAttributes(VSAs) to the RADIUS server by assigning the authentication profile to a GlobalProtect portal or gateway. RADIUS administrators can then perform administrative tasks based on those VSAs.

- **Administrative accounts with authorization managed locally on the firewall or Panorama**—Assign the authentication profile to firewall administrator or Panorama administrator accounts.
- **Administrative accounts with authorization managed on the RADIUS server**—The following procedure describes how to configure RADIUS authentication and authorization for firewall administrators. For Panorama administrators, refer to Configure RADIUS Authentication for PanoramaAdministrators.

STEP 1 | Add a RADIUS server profile.

The profile defines how the firewall connects to the RADIUS server.

1. Select **Device** > **Server Profiles** > **RADIUS** and **Add** a profile.
2. Enter a **Profile Name** to identify the server profile.
3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–120).



If you use the server profile to integrate the firewall with an MFA service, enter an interval that gives users enough time to authenticate. For example, if the MFA service prompts for a one-time password (OTP), users need time to see the OTP on their endpoint device and then enter the OTP in the MFA login page.

4. Select the **Authentication Protocol** (default is **CHAP**) that the firewall uses to authenticate to the RADIUS server.



Select CHAP if the RADIUS server supports that protocol; it is more secure than PAP.

5. **Add** each RADIUS server and enter the following:
   - **Name** to identify the server
   - **RADIUS Server** IP address or FQDN. If you use an FQDN to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
   - **Secret**/**Confirm Secret** is a key to encrypt passwords and can be up to 64 characters in length.
   - Server **Port** for authentication requests (default is 1812)
6. Click **OK** to save the server profile.

© 2018 Palo Alto Networks, Inc.

STEP 2 | Assign the RADIUS server profile to an authentication profile.

The authentication profile defines authentication settings that are common to a set of users.

1. Select **Device** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the authentication profile.
3. Set the **Type** to **RADIUS**.
4. Select the **Server Profile** you configured.
5. Select **Retrieve user group from RADIUS** to collect user group information from VSAs defined on the RADIUS server.

   The firewall matches the group information against the groups you specify in the Allow List of the authentication profile.
6. Select **Advanced** and, in the Allow List, **Add** the users and groups that are allowed to authenticate with this authentication profile.
7. Click **OK** to save the authentication profile.

STEP 3 | Configure the firewall to use the authentication profile for all administrators.

1. Select **Device** > **Setup** > **Management** and edit the Authentication Settings.
2. Select the **Authentication Profile** you configured and click **OK**.

STEP 4 | Configure the roles and access domains that define authorization settings for administrators.

If you already defined RADIUS VSAs on the RADIUS server, the names you specify for roles and access domains on the firewall must match the VSA values.

1. Configure an Admin Role Profile if the administrator uses a custom role instead of a predefined (dynamic) role.
2. Configure an access domain if the firewall has more than one virtual system:
   1. Select **Device** > **Access Domain**, **Add** an access domain, and enter a **Name** to identify the access domain.
   2. **Add** each virtual system that the administrator will access, and then click **OK**.

STEP 5 | **Commit** your changes to activate them on the firewall.

STEP 6 | Configure the RADIUS server to authenticate and authorize administrators.

Refer to your RADIUS server documentation for the specific instructions to perform these steps:

1. Add the firewall IP address or hostname as the RADIUS client.
2. Add the administrator accounts.

   If the RADIUS server profile specifies CHAP as the Authentication Protocol, you must define accounts with reversibly encrypted passwords. Otherwise, CHAP authentication will fail.
3. Define the vendor code for the firewall (25461) and define the RADIUS VSAs for the role, access domain, and user group of each administrator.

   For detailed instructions, refer to the following Knowledge Base articles:

   - For Windows 2003 Server, Windows 2008 (and later), and Cisco Secure Access Control Server (ACS) 4.0—RADIUS Vendor-Specific Attributes (VSAs)
   - For Cisco ACS 5.2—Configuring Cisco ACS 5.2 for use with Palo AltoVSAs

   *When configuring the advanced vendor options on the ACS, you must set both the Vendor Length Field Size and Vendor Type Field Size to **1**. Otherwise, authentication will fail.*

© 2018 Palo Alto Networks, Inc.

STEP 7 | Verify that the RADIUS server performs authentication and authorization for administrators.

1. Log in the firewall web interface using an administrator account that you added to the RADIUS server.
2. Verify that you can access only the web interface pages that are allowed for the role you associated with the administrator.
3. In the **Monitor**, **Policies**, and **Objects** tabs, verify that you can access only the virtual systems that are allowed for the access domain you associated with the administrator.

© 2018 Palo Alto Networks, Inc.

Configure LDAP Authentication

You can use LDAP to authenticate end users who access applications or services through Captive Portal and authenticate firewall or Panorama administrators who access the web interface.



You can also connect to an LDAP server to define policy rules based on user groups. For details, see Map Users to Groups.

STEP 1 | Add an LDAP server profile.

The profile defines how the firewall connects to the LDAP server.

1. Select **Device** > **Server Profiles** > **LDAP** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** the LDAP servers (up to four). For each server, enter a **Name** (to identify the server), **LDAP Server** IP address or FQDN, and server **Port** (default 389).

   If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
4. Select the server **Type**.
5. Enter the **Bind Timeout** and **Search Timeout** in seconds (default is 30 for both).
6. Click **OK** to save the server profile.

STEP 2 | Assign the server profile to an Configure an Authentication Profile and Sequence to define various authentication settings.

STEP 3 | Assign the authentication profile to the firewall application that requires authentication.

- **Administrative access to the web interface**—Configure a Firewall Administrator Account and assign the authentication profile you configured.
- **End user access to services and applications**—For the full procedure to configure authentication for end users, see Configure Authentication Policy.

STEP 4 | Verify that the firewall can Test Authentication Server Connectivity to authenticate users.

© 2018 Palo Alto Networks, Inc.

Connection Timeouts for Authentication Servers

You can configure the firewall to use External Authentication Services for authenticating administrators who access the firewall or Panorama and end users who access services or applications through Captive Portal. To ensure that the firewall does not waste resources by continuously trying to reach an authentication server that is unreachable, you can set a timeout interval after which the firewall stops trying to connect. You set the timeout in the server profiles that define how the firewall connects to the authentication servers. When choosing timeout values, your goal is to strike a balance between the need to conserve firewall resources and to account for normal network delays that affect how quickly authentication servers respond to the firewall.

- Guidelines for Setting Authentication Server Timeouts
- Modify the PAN-OS Web Server Timeout
- Modify the Captive Portal Session Timeout

Guidelines for Setting Authentication Server Timeouts

The following are some guidelines for setting the timeouts for firewall attempts to connect with External Authentication Services.

- In addition to the timeouts you set in server profiles for specific servers, the firewall has a global PAN-OS web server timeout. This global timeout applies when the firewall connects to any external server for authenticating administrative access to the firewall web interface or PAN-OS XML API and end user access to applications or services through Captive Portal. The global timeout is 30 seconds by default (range is 3 to 125). It must be the same as or greater than the total time that any server profile allows for connection attempts. The total time in a server profile is the timeout value multiplied by the number of retries and the number of servers. For example, if a RADIUS server profile specifies a 3-second timeout, 3 retries, and 4 servers, the total time that the profile allows for connection attempts is 36 seconds (3 x 3 x 4). Modify the PAN-OS Web Server Timeout if necessary.

  Do not change the PAN-OS web server timeout unless you see authentication failures. Setting the timeout too high could degrade the performance of the firewall or cause it to drop authentication requests. You can review authentication failures in Authentication logs.

- The firewall applies a Captive Portal session timeout that defines how long end users can take to respond to the authentication challenge in a Captive Portal web form. The web form displays when users request services or applications that match an Authentication policy rule. The session timeout is 30 seconds by default (range is 1 to 1,599,999). It must be the same as or greater than the PAN-OS web server timeout. Modify the Captive Portal Session Timeout if necessary. Keep in mind that increasing the PAN-OS web server and Captive Portal session timeouts might degrade the performance of the firewall or cause it to drop authentication requests.

  The Captive Portal session timeout is not related to the timers that determine how long the firewall retains IP address-to-username mappings.

- Timeouts are cumulative for authentication sequences. For example, consider the case of an authentication sequence with two authentication profiles. One authentication profile specifies a RADIUS server profile with a 3-second timeout, 3 retries, and 4 servers. The other authentication profile specifies a TACACS+ server profile with a 3-second timeout and 2 servers. The longest possible period in which the firewall can try to authenticate user accounts with that authentication sequence is 42 seconds: 36 seconds for the RADIUS server profile plus 6 seconds for the TACACS+ server profile.

© 2018 Palo Alto Networks, Inc.

- The non-configurable timeout for Kerberos servers is 17 seconds for each server specified in the Kerberos server profile.
- To configure the timeouts and related settings for other server types, see:
  - Add an MFA server profile.
  - Add a SAML IdP server profile.
  - Add a TACACS+ server profile.
  - Add a RADIUS server profile.
  - Add an LDAP server profile.

Modify the PAN-OS Web Server Timeout

The PAN-OS web server timeout must be the same as or greater than the timeout in any authentication server profile multiplied by the number of retries and the number of servers in that profile.



Do not change the PAN-OS web server timeout unless you see authentication failures. Setting the timeout too high could degrade the performance of the firewall or cause it to drop authentication requests. You can review authentication failures in Authentication logs.

STEP 1 | Access the firewall CLI.

STEP 2 | Set the PAN-OS web server timeout by entering the following commands, where ***<value>*** is the number of seconds (default is 30; range is 3 to 125).

```
> configure
# set deviceconfig setting l3-service timeout <value>
# commit
```

Modify the Captive Portal Session Timeout

The Captive Portal session timeout must be the same as or greater than the PAN-OS web server timeout. For details, see Connection Timeouts for Authentication Servers.



The more you raise the PAN-OS web server and Captive Portal session timeouts, the slower Captive Portal will respond to users.

STEP 1 | Select **Device** > **Setup** > **Session** and edit the Session Timeouts.

STEP 2 | Enter a new **Captive Portal** value in seconds (default is 30; range is 1 to 1,599,999) and click **OK**.

STEP 3 | **Commit** your changes.

© 2018 Palo Alto Networks, Inc.

Configure Local Database Authentication

You can configure a user database that is local to the firewall to authenticate administrators who access the firewall web interface and to authenticate end users who access applications through Captive Portal or GlobalProtect. Perform the following steps to configure Local Authentication with a local database.



External Authentication Services are usually preferable to local authentication because they provide the benefit of central account management.

You can also configure local authentication without a database, but only for firewall or Panorama administrators.

STEP 1 | Add the user account to the local database.

1. Select **Device** > **Local User Database** > **Users** and click **Add**.
2. Enter a user **Name** for the administrator.
3. Enter a **Password** and **Confirm Password** or enter a **Password Hash**.
4. **Enable** the account (enabled by default) and click **OK**.

STEP 2 | Add the user group to the local database.

Required if your users require group membership.

1. Select **Device** > **Local User Database** > **User Groups** and click **Add**.
2. Enter a **Name** to identify the group.
3. **Add** each user who is a member of the group and click **OK**.

STEP 3 | Configure an authentication profile.

The authentication profile defines authentication settings that are common to a set of users. Set the authentication **Type** to **Local Database**.

STEP 4 | Assign the authentication profile to an administrator account or to an Authentication policy rule for end users.

- **Administrators**—Configure a Firewall Administrator Account:

  Specify the **Name** of a user you defined earlier in this procedure.

  Assign the **Authentication Profile** that you configured for the account.

- **End users**—For the full procedure to configure authentication for end users, see Configure Authentication Policy.

STEP 5 | Verify that the firewall can Test Authentication Server Connectivity to authenticate users.

© 2018 Palo Alto Networks, Inc.

Configure an Authentication Profile and Sequence

An authentication profile defines the authentication service that validates the login credentials of administrators who access the firewall web interface and end users who access applications through Captive Portal or GlobalProtect. The service can be Local Authentication that the firewall provides or External Authentication Services. The authentication profile also defines options such as Kerberos single sign-on (SSO).

Some networks have multiple databases (such as TACACS+ and LDAP) for different users and user groups. To authenticate users in such cases, configure an *authentication sequence*—a ranked order of authentication profiles that the firewall matches a user against during login. The firewall checks against each profile in sequence until one successfully authenticates the user. If the sequence includes an authentication profile that specifies local database authentication, the firewall always checks that profile first regardless of the order in the sequence. A user is denied access only if authentication fails for all the profiles in the sequence. The sequence can specify authentication profiles that are based on any authentication service that the firewall supports excepts Multi-Factor Authentication (MFA) and SAML.

STEP 1 | (External service only) Enable the firewall to connect to an external server for authenticating users:

1. Set up the external server. Refer to your server documentation for instructions.
2. Configure a server profile for the type of authentication service you use.
   - Add a RADIUS server profile.

     If the firewall integrates with an MFA service through RADIUS, you must add a RADIUS server profile. In this case, the MFA service provides all the authentication factors. If the firewall integrates with an MFA service through a vendor API, you can still use a RADIUS server profile for the first factor but MFA server profiles are required for additional factors.

   - Add an MFA server profile.
   - Add a SAML IdP server profile.
   - Add a Kerberos server profile.
   - Add a TACACS+ server profile.
   - Add an LDAP server profile.

STEP 2 | (Local database authentication only) Configure a user database that is local to the firewall.

Perform these steps for each user and user group for which you want to configure Local Authentication based on a user identity store that is local to the firewall:

1. Add the user account to the local database.
2. (Optional) Add the user group to the local database.

STEP 3 | (Kerberos SSO only) Create a Kerberos keytab for the firewall if Kerberos single sign-on (SSO) is the primary authentication service.

Create a Kerberos keytab. A keytab is a file that contains Kerberos account information for the firewall. To support Kerberos SSO, your network must have a Kerberos infrastructure.

STEP 4 | Configure an authentication profile.

© 2018 Palo Alto Networks, Inc.

Define one or both of the following:

- **Kerberos SSO**—The firewall first tries SSO authentication. If that fails, it falls back to the specified authentication **Type**.
- **External authentication or local database authentication**—The firewall prompts the user to enter login credentials, and uses an external service or local database to authenticate the user.

1. Select **Device** > **Authentication Profile** and **Add** the authentication profile.
2. Enter a **Name** to identify the authentication profile.
3. Select the **Type** of authentication service.

   If you use Multi-Factor Authentication, the selected type applies only to the first authentication factor. You select services for additional MFA factors in the **Factors** tab.

   If you select **RADIUS**, **TACACS+**, **LDAP**, or **Kerberos**, select the **Server Profile**.

   If you select **LDAP**, select the **Server Profile** and define the **Login Attribute**. For Active Directory, enter **sAMAccountName** as the value.

   If you select **SAML**, select the **IdP Server Profile**.
4. If you want to enable Kerberos SSO, enter the **Kerberos Realm** (usually the DNS domain of the users, except that the realm is UPPERCASE) and **Import** the **Kerberos Keytab** that you created for the firewall or Panorama.
5. (MFA only) Select **Factors**, **Enable Additional Authentication Factors**, and **Add** the MFA server profiles you configured.

   The firewall will invoke each MFA service in the listed order, from top to bottom.
6. Select **Advanced** and **Add** the users and groups that can authenticate with this profile.

   You can select users and groups from the local database or, if you configured the firewall to Map Users to Groups, from an LDAP-based directory service such as Active Directory. By default, the list is empty, meaning no users can authenticate.

   You can also select custom groups defined in a group mapping configuration.
7. Click **OK** to save the authentication profile.

STEP 5 | Configure an authentication sequence.

Required if you want the firewall to try multiple authentication profiles to authenticate users. The firewall evaluates the profiles in top-to-bottom order until one profile successfully authenticates the user.

1. Select **Device** > **Authentication Sequence** and **Add** the authentication sequence.
2. Enter a **Name** to identify the authentication sequence.

   To expedite the authentication process, Use domain to determine authentication profile: the firewall matches the domain name that a user enters during login with the User Domain or Kerberos Realm of an authentication profile in the sequence, and then uses that profile to authenticate the user. If the firewall does not find a match, or if you disable the option, the firewall tries the profiles in the top-to-bottom sequence.
3. **Add** each authentication profile. To change the evaluation order of the profiles, select a profile and **Move Up** or **Move Down**.
4. Click **OK** to save the authentication sequence.

STEP 6 | Assign the authentication profile or sequence to an administrative account for firewall administrators or to Authentication policy for end users.

© 2018 Palo Alto Networks, Inc.

- **Administrators**—Assign the authentication profile based on how you manager administrator authorization:

  Authorization managed locally on the firewall—Configure a Firewall Administrator Account.

  Authorization managed on a SAML, TACACS+, or RADIUS server—Select **Device** > **Setup** > **Management**, edit the Authentication Settings, and select the **Authentication Profile**.

- **End users**—For the full procedure to configure authentication for end users, see Configure Authentication Policy.

STEP 7 | Verify that the firewall can Test Authentication Server Connectivity to authenticate users.

© 2018 Palo Alto Networks, Inc.

Test Authentication Server Connectivity

The test authentication feature enables you to verify whether the firewall or Panorama can communicate with the authentication server specified in an authentication profile and whether an authentication request succeeds for a specific user. You can test authentication profiles that authenticate administrators who access the web interface or that authenticate end users who access applications through GlobalProtect or Captive Portal. You can perform authentication tests on the candidate configuration to verify the configuration is correct before committing.

STEP 1 | Configure an authentication profile. You do not need to commit the authentication profile or server profile configuration before testing.

STEP 2 | Log into the firewall CLI.

STEP 3 | (Firewalls with multiple virtual systems) Define the target virtual system that the test command will access.

This is required on firewalls with multiple virtual systems so that the test authentication command can locate the user you will test.

Define the target virtual system by entering:

```
admin@PA-3060> set system setting target-vsys <vsys-name>
```

For example, if the user is defined in vsys2, enter:

```
admin@PA-3060> set system setting target-vsys vsys2
```

*The **target-vsys** option is per login session; the firewall clears the option when you log off.*

STEP 4 | Test the authentication profile by entering the following command:

```
admin@PA-3060> test authentication authentication-profile <authentication-
profile-name> username <username> password
```

For example, to test an authentication profile named **my-profile** for a user named **bsimpson**, enter:

```
admin@PA-3060> test authentication authentication-profile my-profile
 username bsimpson password
```

*When running the **test** command, the names of authentication profiles and server profiles are case sensitive. Also, if an authentication profile has a username modifier defined, you must enter the modifier with the username. For example, if you add the username modifier **%USERINPUT%@%USERDOMAIN%** for a user named **bsimpson** and the domain name is **mydomain.com**, enter **bsimpson@mydomain.com** as the username. This ensures that the firewall sends the correct credentials to the authentication server. In this example, mydomain.com is the domain that you define in the User Domain field in the authentication profile.*

STEP 5 | View the test output.

© 2018 Palo Alto Networks, Inc.

If the authentication profile is configured correctly, the output displays `Authentication succeeded`. If there is a configuration issue, the output displays information to help you troubleshoot the configuration.



The output results vary based on several factors related to the authentication type that you are testing as well as the type of issue. For example, RADIUS and TACACS+ use different underlying libraries, so the same issue that exists for both of these types will produce different errors. Also, if there is a network problem, such as using an incorrect port or IP address in the authentication server profile, the output error is not specific. This is because the test command cannot perform the initial handshake between the firewall and the authentication server to determine details about the issue.

© 2018 Palo Alto Networks, Inc.

Authentication Policy

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 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 One-time Password (OTP)authentication. For the first factor, users authenticate through a Captive Portal web form. For any additional factors, users authenticate through a Multi-Factor Authentication (MFA) login page.



To implement Authentication policy for GlobalProtect, refer to Configure GlobalProtect to Facilitate Multi-Factor Authentication Notifications*.*

After the user authenticates for all factors, the firewall evaluates Security Policy to determine whether 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. 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 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 User-ID or Authentication logs and forward the logs to Panorama or external services as you would for any other log types.

- Authentication Timestamps
- Configure Authentication Policy

Authentication Timestamps

When configuring an Authentication policy rule, you can specify a timeout period during which a user authenticates only for initial access to services and applications, not for subsequent access. Your goal is to specify a timeout that strikes a balance between the need to secure services and applications and the need to minimize interruptions to the user workflow. When a user authenticates, the firewall records a timestamp for the first authentication challenge (factor) and a timestamp for any additional Multi-Factor Authentication (MFA) factors. When the user subsequently requests services and applications that match an Authentication rule, the firewall evaluates the timeout specified in the rule relative to each timestamp. This means the firewall reissues authentication challenges on a per-factor basis when timeouts expire. If you Redistribute User Mappings and Authentication Timestamps, all your firewalls will enforce Authentication policy timeouts consistently for all users.



The firewall records a separate timestamp for each MFA vendor. For example, if you use Duo v2 and PingID servers to issue challenges for MFA factors, the firewall records one timestamp for the response to the Duo factor and one timestamp for the response to the PingID factor.

Within the timeout period, a user who successfully authenticates for one Authentication rule can access services or applications that other rules protect. However, this portability applies only to rules that trigger the same authentication factors. For example, a user who successfully authenticates for a rule that triggers TACACS+ authentication must authenticate again for a rule that triggers SAML authentication, even if the access requests are within the timeout period for both rules.

© 2018 Palo Alto Networks, Inc.

When evaluating the timeout in each Authentication rule and the global timer defined in the Captive Portal settings (see Configure Captive Portal), the firewall prompts the user to re-authenticate for whichever setting expires first. Upon re-authenticating, the firewall records new authentication timestamps for the rules and resets the time count for the Captive Portal timer. Therefore, to enable different timeout periods for different Authentication rules, set the Captive Portal timer to a value that is the same as or higher than the timeout in any rule.

Configure 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.

STEP 1 | Configure Captive Portal. If you use Multi-Factor Authentication (MFA) services to authenticate users, you must set the **Mode** to **Redirect**.

STEP 2 | 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 single sign-on (SSO)—Create a Kerberos keytab for the firewall. Optionally, you can configure the firewall to use Kerberos SSO as the primary authentication service and, if SSO failures occur, fall back to an external service or local database authentication.

STEP 3 | Configure an Authentication Profile and 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**.

STEP 4 | 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.

1. Select **Objects** > **Authentication** and **Add** an object.
2. Enter a **Name** to identify the object.
3. 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 to the first authentication factor instead of having the user enter login credentials. For this method, you must have configured Kerberos SSO in the authentication profile or NT LAN Manager (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.
4. Select the **Authentication Profile** you configured.
5. Enter the **Message** that the Captive Portal web form will display to tell users how to authenticate for the first authentication factor.
6. Click **OK** to save the object.

© 2018 Palo Alto Networks, Inc.

STEP 5 | Configure an Authentication policy rule.

Create a rule for each set of users, services, and URL categories that require the same authentication services and settings.

1. Select **Policies** > **Authentication** and **Add** a rule.
2. Enter a **Name** to identify the rule.
3. Select **Source** and **Add** specific zones and IP addresses or select **Any** zones or IP addresses.

   The rule applies only to traffic coming from the specified IP addresses or from interfaces in the specified zones.
4. Select **User** and select or **Add** the source users and user groups to which the rule applies (default is **any**).
5. Select or **Add** the Host Information Profiles to which the rule applies (default is **any**).
6. Select **Destination** and **Add** specific zones and IP addresses or select **any** zones or IP addresses.

   The IP addresses can be resources (such as servers) for which you want to control access.
7. Select **Service/URL Category** and select or **Add** the services and service groups for which the rule controls access (default is **service-http**).
8. 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 your most sensitive internal sites.
9. Select **Actions** and select the **Authentication Enforcement** object you created.
10. Specify the **Timeout** period in minutes (default 60) during which the firewall prompts the user to authenticate only once for repeated access to services and applications.
11. Click **OK** to save the rule.

STEP 6 | (MFA only) Customize the MFA login page.

The firewall displays this page so that users can authenticate for any additional MFA factors.

STEP 7 | Verify that the firewall enforces Authentication policy.

1. Log in to your network as one of the source users specified in an Authentication policy rule.
2. Request a service or URL category that matches one specified in the rule.

   The firewall displays the Captive Portal web form for the first authentication factor. For example:



If you configured the firewall to use one or more MFA services, authenticate for the additional authentication factors.

3. End the session for the service or URL you just accessed.
4. Start a new session for the same service or application. Be sure to perform this step within the **Timeout** period you configured in the Authentication rule.

   The firewall allows access without re-authenticating.
5. Wait until the **Timeout** period expires and request the same service or application.

   The firewall prompts you to re-authenticate.

© 2018 Palo Alto Networks, Inc.

STEP 8 | (Optional) Redistribute User Mappings and Authentication Timestamps to other firewalls that enforce Authentication policy to ensure they all apply the timeouts consistently for all users.

© 2018 Palo Alto Networks, Inc.

Troubleshoot Authentication Issues

When users fail to authenticate to a Palo Alto Networks firewall or Panorama, or the Authentication process takes longer than expected, analyzing authentication-related information can help you determine whether the failure or delay resulted from:

- **User behavior**—For example, users are locked out after entering the wrong credentials or a high volume of users are simultaneously attempting access.
- **System or network issues**—For example, an authentication server is inaccessible.
- **Configuration issues**—For example, the Allow List of an authentication profile doesn't have all the users it should have.

The following CLI commands display information that can help you troubleshoot these issues:

Task	Command
Display the number of locked user accounts associated with the authentication profile (**auth-profile**), authentication sequence (**is-seq**), or virtual system (**vsys**).  *To unlock users, use the following operational command:* **> request authentication [unlock-admin | unlock-user]**	**show authentication locked-users** **{** **vsys** ***<value>*** **|** **auth-profile** ***<value>*** **|** **is-seq** **{yes | no}** **{auth-profile | vsys}** ***<value>*** **}**
Use the **debug authentication** command to troubleshoot authentication events. Use the **show** options to display authentication request statistics and the current debugging level: • **show** displays the current debugging level for the authentication service (authd). • **show-active-requests** displays the number of active checks for authentication requests, allow lists, locked user accounts, and Multi-Factor Authentication (MFA) requests. • **show-pending-requests** displays the number of pending checks for authentication requests, allow lists, locked user accounts, and MFA requests. • **connection-show** displays authentication request and response statistics for all authentication servers or for a specific protocol type. Use the **connection-debug** options to enable or disable authentication debugging:	**debug authentication** **{** **on {debug | dump | error | info | warn} |** **show |** **show-active-requests |** **show-pending-requests |** **connection-show |** **{** **connection-id |** **protocol-type** **{** **Kerberos connection-id** ***<value>*** **|** **LDAP connection-id** ***<value>*** **|** **RADIUS connection-id** ***<value>*** **|** **TACACS+ connection-id** ***<value>*** **|** **}** **connection-debug-on |** **{** **connection-id |** **debug-prefix |**

© 2018 Palo Alto Networks, Inc.

Task	Command
• Use the **on** option to enable or the **off** option to disable debugging for authd. • Use the **connection-debug-on** option to enable or the **connection-debug-off** option to disable debugging for all authentication servers or for a specific protocol type.	**protocol-type** **{** **Kerberos connection-id** ***<value>*** **|** **LDAP connection-id** ***<value>*** **|** **RADIUS connection-id** ***<value>*** **|** **TACACS+ connection-id** ***<value>*** **|** **}** **connection-debug-off |** **{** **connection-id |** **protocol-type** **{** **Kerberos connection-id** ***<value>*** **|** **LDAP connection-id <value> |** **RADIUS connection-id <value> |** **TACACS+ connection-id <value> |** **}** **connection-debug-on** **}**

Certificate Management

The following topics describe the different keys and certificates that Palo Alto Networks® firewalls and Panorama use, and how to obtain and manage them:

- Keys and Certificates
- Certificate Revocation
- Certificate Deployment
- Set Up Verification for Certificate Revocation Status
- Configure the Master Key
- Obtain Certificates
- Export a Certificate and Private Key
- Configure a Certificate Profile
- Configure an SSL/TLS Service Profile
- Replace the Certificate for Inbound Management Traffic
- Configure the Key Size for SSL Forward Proxy Server Certificates
- Revoke and Renew Certificates
- Secure Keys with a Hardware Security Module

© 2018 Palo Alto Networks, Inc.

Keys and Certificates

To ensure trust between parties in a secure communication session, Palo Alto Networks firewalls and Panorama use digital certificates. Each certificate contains a cryptographic key to encrypt plaintext or decrypt ciphertext. Each certificate also includes a digital signature to authenticate the identity of the issuer. The issuer must be in the list of trusted certificate authorities (CAs) of the authenticating party. Optionally, the authenticating party verifies the issuer did not revoke the certificate (see Certificate Revocation).

Palo Alto Networks firewalls and Panorama use certificates in the following applications:

- User authentication for Captive Portal, GlobalProtect™, Mobile Security Manager, and web interface access to a firewall or Panorama.
- Device authentication for GlobalProtect VPN (remote user-to-site or large scale).
- Device authentication for IPSec site-to-site VPN with Internet Key Exchange (IKE).
- Decrypting inbound and outbound SSL traffic.

  A firewall decrypts the traffic to apply policy rules, then re-encrypts it before forwarding the traffic to the final destination. For outbound traffic, the firewall acts as a forward proxy server, establishing an SSL/TLS connection to the destination server. To secure a connection between itself and the client, the firewall uses a *signing certificate* to automatically generate a copy of the destination server certificate.

The following table describes the keys and certificates that Palo Alto Networks firewalls and Panorama use. As a best practice, use different keys and certificates for each usage.

Table 1: Palo Alto Networks Device Keys/Certificates

Key/Certificate Usage	Description
Administrative Access	Secure access to firewall or Panorama administration interfaces (HTTPS access to the web interface) requires a server certificate for the MGT interface (or a designated interface on the dataplane if the firewall or Panorama does not use MGT) and, optionally, a certificate to authenticate the administrator.
Captive Portal	In deployments where Authentication policy identifies users who access HTTPS resources, designate a server certificate for the Captive Portal interface. If you configure Captive Portal to use certificates for identifying users (instead of, or in addition to, interactive authentication), deploy client certificates also. For more information on Captive Portal, see Map IP Addresses to Usernames Using Captive Portal.
Forward Trust	For outbound SSL/TLS traffic, if a firewall acting as a forward proxy trusts the CA that signed the certificate of the destination server, the firewall uses the forward trust CA certificate to generate a copy of the destination server certificate to present to the client. To set the private key size, see Configure the Key Size for SSL Forward Proxy Server Certificates. For added security, store the key on a hardware security module (for details, see Secure Keys with a Hardware Security Module).
Forward Untrust	For outbound SSL/TLS traffic, if a firewall acting as a forward proxy does not trust the CA that signed the certificate of the destination server, the firewall uses the forward untrust CA certificate to generate a copy of the destination server certificate to present to the client.

© 2018 Palo Alto Networks, Inc.

Key/Certificate Usage	Description
SSL Inbound Inspection	The keys that decrypt inbound SSL/TLS traffic for inspection and policy enforcement. For this application, import onto the firewall a private key for each server that is subject to SSL/TLS inbound inspection. See Configure SSL Inbound Inspection. *Beginning in PAN-OS 8.0, firewalls use the Elliptic-Curve Diffie-Hellman Ephemeral (ECDHE) algorithm to perform strict certificate checking. This means that if the firewall uses an intermediate certificate, you must reimport the certificate from your web server to the firewall after you upgrade to a PAN-OS 8.0 or later release and combine the server certificate with the intermediate certificate (install a chained certificate). Otherwise, SSL Inbound Inspection sessions that have an intermediate certificate in the chain will fail. To install a chained certificate:* ***1.*** *Open each certificate (.cer) file in a plain-text editor such as Notepad.* ***2.*** *Paste each certificate end-to-end with the Server Certificate at the top with each signer included below.* ***3.*** *Save the file as a text (.txt) or certificate (.cer) file (the name of the file cannot contain blank spaces).* ***4.*** *Import the combined (chained) certificate into the firewall.*
SSL Exclude Certificate	Certificates for servers to exclude from SSL/TLS decryption. For example, if you enable SSL decryption but your network includes servers for which the firewall should not decrypt traffic (for example, web services for your HR systems), import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Decryption Exclusions.
GlobalProtect	All interaction among GlobalProtect components occurs over SSL/TLS connections. Therefore, as part of the GlobalProtect deployment, deploy server certificates for all GlobalProtect portals, gateways, and Mobile Security Managers. Optionally, deploy certificates for authenticating users also. *The GlobalProtect Large Scale VPN (LSVPN) feature requires a CA signing certificate.*
Site-to-Site VPNs (IKE)	In a site-to-site IPSec VPN deployment, peer devices use Internet Key Exchange (IKE) gateways to establish a secure channel. IKE gateways use certificates or preshared keys to authenticate the peers to each other. You configure and assign the certificates or keys when defining an IKE gateway on a firewall. See Site-to-Site VPN Overview.
Master Key	The firewall uses a master key to encrypt all private keys and passwords. If your network requires a secure location for storing private keys, you can use an encryption (wrapping) key stored on a hardware security module (HSM) to encrypt the master key. For details, see Encrypt a Master Key Using an HSM.
Secure Syslog	The certificate to enable secure connections between the firewall and a syslog server. See Syslog Field Descriptions.

© 2018 Palo Alto Networks, Inc.

Key/Certificate Usage	Description
Trusted Root CA	The designation for a root certificate issued by a CA that the firewall trusts. The firewall can use a self-signed root CA certificate to automatically issue certificates for other applications (for example, SSL Forward Proxy). Also, if a firewall must establish secure connections with other firewalls, the root CA that issues their certificates must be in the list of trusted root CAs on the firewall.
Inter-Device Communication	By default, Panorama, firewalls, and Log Collectors use a set of predefined certificates for the SSL/TLS connections used for management and log forwarding. However, you can enhance these connection by deploying custom certificates to the devices in your deployment. These certificates can also be used to secure the SSL/TLS connection between Panorama HA peers.

© 2018 Palo Alto Networks, Inc.

Certificate Revocation

Palo Alto Networks firewalls and Panorama use digital certificates to ensure trust between parties in a secure communication session. Configuring a firewall or Panorama to check the revocation status of certificates provides additional security. A party that presents a revoked certificate is not trustworthy. When a certificate is part of a chain, the firewall or Panorama checks the status of every certificate in the chain except the root CA certificate, for which it cannot verify revocation status.

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority that issued the certificate must revoke it.

The firewall and Panorama support the following methods for verifying certificate revocation status. If you configure both methods, the firewall or Panorama first tries the OCSP method; if the OCSP server is unavailable, it uses the CRL method.

- Certificate Revocation List (CRL)
- Online Certificate Status Protocol (OCSP)



In PAN-OS, certificate revocation status verification is an optional feature. It is a best practice to enable it for certificate profiles, which define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPSec VPN, and web interface access to the firewall or Panorama.

Certificate Revocation List (CRL)

Each certificate authority (CA) periodically issues a certificate revocation list (CRL) to a public repository. The CRL identifies revoked certificates by serial number. After the CA revokes a certificate, the next CRL update will include the serial number of that certificate.

The Palo Alto Networks firewall downloads and caches the last-issued CRL for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the CRL for the issuing CA. Also, the cache only stores a CRL until it expires.

The firewall supports CRLs only in Distinguished Encoding Rules (DER) format. If the firewall downloads a CRL in any other format—for example, Privacy Enhanced Mail (PEM) format—any revocation verification process that uses that CRL will fail when a user performs an activity that triggers the process (for example, sending outbound SSL data). The firewall will generate a system log for the verification failure. If the verification was for an SSL certificate, the firewall will also display the SSL Certificate Errors Notify response page to the user.

To use CRLs for verifying the revocation status of certificates used for the decryption of inbound and outbound SSL/TLS traffic, see Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

To use CRLs for verifying the revocation status of certificates that authenticate users and devices, configure a certificate profile and assign it to the interfaces that are specific to the application: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, or web interface access to Palo Alto Networks firewalls or Panorama. For details, see Configure Revocation Status Verification of Certificates.

© 2018 Palo Alto Networks, Inc.

Online Certificate Status Protocol (OCSP)

When establishing an SSL/TLS session, clients can use Online Certificate Status Protocol (OCSP) to check the revocation status of the authentication certificate. The authenticating client sends a request containing the serial number of the certificate to the OCSP responder (server). The responder searches the database of the certificate authority (CA) that issued the certificate and returns a response containing the status (good, revoked or unknown) to the client. The advantage of the OCSP method is that it can verify status in real-time, instead of depending on the issue frequency (hourly, daily, or weekly) of CRLs.

The Palo Alto Networks firewall downloads and caches OCSP status information for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the OCSP information for the issuing CA. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall as an OCSP responder (see Configure an OCSP Responder).

To use OCSP for verifying the revocation status of certificates when the firewall functions as an SSL forward proxy, perform the steps under Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption.

The following applications use certificates to authenticate users and/or devices: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPSec VPN, and web interface access to Palo Alto Networks firewalls or Panorama. To use OCSP for verifying the revocation status of the certificates:

- Configure an OCSP responder (if you are configuring the firewall as an OCSP responder).
- Enable the HTTP OCSP service on the firewall (if you are configuring the firewall as an OCSP responder).
- Create or obtain a certificate for each application.
- Configure a certificate profile for each application.
- Assign the certificate profile to the relevant application.

To cover situations where the OCSP responder is unavailable, configure CRL as a fall-back method. For details, see Configure Revocation Status Verification of Certificates.

© 2018 Palo Alto Networks, Inc.

Certificate Deployment

The basic approaches to deploy certificates for Palo Alto Networks firewalls or Panorama are:

- **Obtain certificates from a trusted third-party CA**—The benefit of obtaining a certificate from a trusted third-party certificate authority (CA) such as VeriSign or GoDaddy is that end clients will already trust the certificate because common browsers include root CA certificates from well-known CAs in their trusted root certificate stores. Therefore, for applications that require end clients to establish secure connections with the firewall or Panorama, purchase a certificate from a CA that the end clients trust to avoid having to pre-deploy root CA certificates to the end clients. (Some such applications are a GlobalProtect portal or GlobalProtect Mobile Security Manager.) However, most third-party CAs cannot issue signing certificates. Therefore, this type of certificate is not appropriate for applications (for example, SSL/TLS decryption and large-scale VPN) that require the firewall to issue certificates. See Obtain a Certificate from an External CA.
- **Obtain certificates from an enterprise CA**—Enterprises that have their own internal CA can use it to issue certificates for firewall applications and import them onto the firewall. The benefit is that end clients probably already trust the enterprise CA. You can either generate the needed certificates and import them onto the firewall, or generate a certificate signing request (CSR) on the firewall and send it to the enterprise CA for signing. The benefit of this method is that the private key does not leave the firewall. An enterprise CA can also issue a signing certificate, which the firewall uses to automatically generate certificates (for example, for GlobalProtect large-scale VPN or sites requiring SSL/TLS decryption). See Import a Certificate and Private Key.
- **Generate self-signed certificates**—You can Create a Self-Signed Root CA Certificate on the firewall and use it to automatically issue certificates for other firewall applications.



If you use this method to generate certificates for an application that requires an end client to trust the certificate, end users will see a certificate error because the root CA certificate is not in their trusted root certificate store. To prevent this, deploy the self-signed root CA certificate to all end user systems. You can deploy the certificates manually or use a centralized deployment method such as an Active Directory Group Policy Object (GPO).

© 2018 Palo Alto Networks, Inc.

Set Up Verification for Certificate Revocation Status

To verify the revocation status of certificates, the firewall uses Online Certificate Status Protocol (OCSP) and/or certificate revocation lists (CRLs). For details on these methods, see Certificate Revocation If you configure both methods, the firewall first tries OCSP and only falls back to the CRL method if the OCSP responder is unavailable. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall to function as the OCSP responder.

The following topics describe how to configure the firewall to verify certificate revocation status:

- Configure an OCSP Responder
- Configure Revocation Status Verification of Certificates
- Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

Configure an OCSP Responder

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of certificates, you must configure the firewall to access an OCSP responder (server). The entity that manages the OCSP responder can be a third-party certificate authority (CA). If your enterprise has its own public key infrastructure (PKI), you can use external OCSP responders or you can configure the firewall itself as an OCSP responder. For details on OCSP, see Certificate Revocation.

STEP 1 | Define an external OCSP responder or configure the firewall itself as an OCSP responder.

1. Select **Device** > **Certificate Management** > **OCSP Responder** and click **Add**.
2. Enter a **Name** to identify the responder (up to 31 characters). The name is case-sensitive. It must be unique and use only letters, numbers, spaces, hyphens, and underscores.
3. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.
4. In the **Host Name** field, enter the host name (recommended) or IP address of the OCSP responder. You can enter an IPv4 or IPv6 address. From this value, PAN-OS automatically derives a URL and adds it to the certificate being verified.

   If you configure the firewall itself as an OCSP responder, the host name must resolve to an IP address in the interface that the firewall uses for OCSP services.
5. Click **OK**.

STEP 2 | If you want the firewall to use the management interface for the OCSP responder interface, enable OCSP communication on the firewall. Otherwise, continue to the next step to configure an alternate interface.

1. Select **Device** > **Setup** > **Management**.
2. In the Management Interface Settings section, edit to select the **HTTP OCSP** check box, then click **OK**.

STEP 3 | **(Optional)** To use an alternate interface as the OCSP responder interface, add an Interface Management Profile to the interface used for OCSP services.

1. Select **Network** > **Network Profiles** > **Interface Mgmt**.
2. Click **Add** to create a new profile or click the name of an existing profile.
3. Select the **HTTP OCSP** check box and click **OK**.

© 2018 Palo Alto Networks, Inc.

4. Select **Network** > **Interfaces** and click the name of the interface that the firewall will use for OCSP services. The OCSP **Host Name** specified in Step Define an OCSP responder. must resolve to an IP address in this interface.
5. Select **Advanced** > **Other info** and select the Interface Management Profile you configured.
6. Click **OK** and **Commit**.

Configure Revocation Status Verification of Certificates

The firewall and Panorama use certificates to authenticate users and devices for such applications as Captive Portal, GlobalProtect, site-to-site IPSec VPN, and web interface access to the firewall/Panorama. To improve security, it is a best practice to configure the firewall or Panorama to verify the revocation status of certificates that it uses for device/user authentication.

STEP 1 | Configure a Certificate Profile for each application.

Assign one or more root CA certificates to the profile and select how the firewall verifies certificate revocation status.

For details on the certificates that various applications use, see Keys and Certificates

STEP 2 | Assign the certificate profiles to the relevant applications.

The steps to assign a certificate profile depend on the application that requires it.

Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

The firewall decrypts inbound and outbound SSL/TLS traffic to apply security rules and rules, then re-encrypts the traffic before forwarding it. (For details, see SSL Inbound Inspection and SSL Forward Proxy.) You can configure the firewall to verify the revocation status of certificates used for decryption as follows.

Enabling revocation status verification for SSL/TLS decryption certificates will add time to the process of establishing the session. The first attempt to access a site might fail if the verification does not finish before the session times out. For these reasons, verification is disabled by default.

STEP 1 | Define the service-specific timeout intervals for revocation status requests.

1. Select **Device** > **Setup** > **Session** and, in the Session Features section, select **Decryption Certificate Revocation Settings**.
2. Perform one or both of the following steps, depending on whether the firewall will use Online Certificate Status Protocol (OCSP) or the Certificate Revocation List (CRL) method to verify the revocation status of certificates. If the firewall will use both, it first tries OCSP; if the OCSP responder is unavailable, the firewall then tries the CRL method.
   - In the CRL section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the CRL service.
   - In the OCSP section, select the **Enable** check box and enter the **Receive Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the OCSP responder.

   Depending on the **Certificate Status Timeout** value you specify in Step 2, the firewall might register a timeout before either or both of the **Receive Timeout** intervals pass.

STEP 2 | Define the total timeout interval for revocation status requests.

© 2018 Palo Alto Networks, Inc.

Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies the session-blocking logic you optionally define in Step 3. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:

- If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.
- If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.
- If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value.

STEP 3 | Define the blocking behavior for unknown certificate status or a revocation status request timeout.

If you want the firewall to block SSL/TLS sessions when the OCSP or CRL service returns a certificate revocation status of unknown, select the **Block Session With Unknown Certificate Status** check box. Otherwise, the firewall proceeds with the session.

If you want the firewall to block SSL/TLS sessions after it registers a request timeout, select the **Block Session On Certificate Status Check Timeout** check box. Otherwise, the firewall proceeds with the session.

STEP 4 | Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Configure the Master Key

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).

If a high availability (HA) configuration, you must use the same master key on both firewalls or Panorama in the pair. Otherwise, HA synchronization will not work properly.

Additionally, if you are using Panorama to manage your firewalls, you must use the same master key on Panorama and all managed firewalls so that Panorama can push configurations to the firewalls.



For added security, Encrypt a Master Key Using an HSM.

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.

STEP 1 | Select **Device** > **Master Key and Diagnostics** and edit the Master Key section.

STEP 2 | Enter the **Current Master Key** if one exists.

STEP 3 | Define a new **New Master Key** and then **Confirm New Master Key**. The key must contain exactly 16 characters.

STEP 4 | To specify the master key **Life Time**, enter the number of **Days** and/or **Hours** after which the key will expire.



You must configure a new master key before the current key expires. If the master key expires, the firewall or Panorama automatically reboots in Maintenance mode. You must then Reset the Firewall to Factory Default Settings.

STEP 5 | Enter a **Time for Reminder** that specifies the number of **Days** and **Hours** before the master key expires when the firewall generates an expiration alarm. The firewall automatically opens the System Alarms dialog to display the alarm.



To ensure the expiration alarm displays, select Device > Log Settings, edit the Alarm Settings, and Enable Alarms.

STEP 6 | (Optional) Select whether to use an **HSM** to encrypt the master key. For details, see Encrypt a Master Key Using an HSM.

STEP 7 | Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Obtain Certificates

- Create a Self-Signed Root CA Certificate
- Generate a Certificate
- Import a Certificate and Private Key
- Obtain a Certificate from an External CA

Create a Self-Signed Root CA Certificate

A self-signed root certificate authority (CA) certificate is the top-most certificate in a certificate chain. A firewall can use this certificate to automatically issue certificates for other uses. For example, the firewall issues certificates for SSL/TLS decryption and for satellites in a GlobalProtect large-scale VPN.

When establishing a secure connection with the firewall, the remote client must trust the root CA that issued the certificate. Otherwise, the client browser will display a warning that the certificate is invalid and might (depending on security settings) block the connection. To prevent this, after generating the self-signed root CA certificate, import it into the client systems.

On a Palo Alto Networks firewall or Panorama, you can generate self-signed certificates only if they are CA certificates.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

STEP 3 | Click **Generate**.

STEP 4 | Enter a **Certificate Name**, such as **`GlobalProtect_CA`**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.

STEP 5 | In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.

STEP 6 | If the firewall has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box.

STEP 7 | Leave the **Signed By** field blank to designate the certificate as self-signed.

STEP 8 | (Required) Select the **Certificate Authority** check box.

STEP 9 | Leave the **OCSP Responder** field blank; revocation status verification doesn't apply to root CA certificates.

STEP 10 | Click **Generate** and **Commit**.

Generate a Certificate

Palo Alto Networks firewalls and Panorama use certificates to authenticate clients, servers, users, and devices in several applications, including SSL/TLS decryption, Captive Portal, GlobalProtect, site-to-site

© 2018 Palo Alto Networks, Inc.

IPSec VPN, and web interface access to the firewall/Panorama. Generate certificates for each usage: for details, see Keys and Certificates.

To generate a certificate, you must first Create a Self-Signed Root CA Certificate or import one (Import a Certificate and Private Key) to sign it. To use Online Certificate Status Protocol (OCSP) for verifying certificate revocation status, Configure an OCSP Responder before generating the certificate.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

STEP 3 | Click **Generate**.

STEP 4 | Select **Local** (default) as the **Certificate Type** unless you want to deploy SCEP certificates to GlobalProtect clients.

STEP 5 | Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.

STEP 6 | In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.

STEP 7 | If the firewall has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box.

STEP 8 | In the **Signed By** field, select the root CA certificate that will issue the certificate.

STEP 9 | (Optional) Select an **OCSP Responder**.

STEP 10 | For the key generation **Algorithm**, select **RSA** (default) or **Elliptical Curve DSA** (ECDSA). ECDSA is recommended for client browsers and operating systems that support it.



Firewalls that run PAN-OS 6.1 and earlier releases will delete any ECDSA certificates that you push from Panorama™, and any RSA certificates signed by an ECDSA certificate authority (CA) will be invalid on those firewalls.

You cannot use a hardware security module (HSM) to store ECDSA keys used for SSL/TLS Decryption.

STEP 11 | Select the **Number of Bits** to define the certificate key length. Higher numbers are more secure but require more processing time.

STEP 12 | Select the **Digest** algorithm. From most to least secure, the options are: **sha512**, **sha384**, **sha256** (default), **sha1**, and **md5**.



Client certificates that are used when requesting firewall services that rely on TLSv1.2 (such as administrator access to the web interface) cannot have sha512 as a digest algorithm. The client certificates must use a lower digest algorithm (such as sha384) or you must limit the Max Version to TLSv1.1 when you Configure an SSL/TLS Service Profile *for the firewall services.*

STEP 13 | For the **Expiration**, enter the number of days (default is 365) for which the certificate is valid.

© 2018 Palo Alto Networks, Inc.

STEP 14 | (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.



If you add a Host Name (DNS name) attribute, it is a best practice for it to match the Common Name. The host name populates the Subject Alternative Name field of the certificate.

STEP 15 | Click **Generate** and, in the Device Certificates page, click the certificate Name.



Regardless of the time zone on the firewall, it always displays the corresponding Greenwich Mean Time (GMT) for certificate validity and expiration dates/times.

STEP 16 | Select the check boxes that correspond to the intended use of the certificate on the firewall.

For example, if the firewall will use this certificate to secure forwarding of syslogs to an external syslog server, select the **Certificate for Secure Syslog** check box.

STEP 17 | Click **OK** and **Commit**.

Import a Certificate and Private Key

If your enterprise has its own public key infrastructure (PKI), you can import a certificate and private key into the firewall from your enterprise certificate authority (CA). Enterprise CA certificates (unlike most certificates purchased from a trusted, third-party CA) can automatically issue CA certificates for applications such as SSL/TLS decryption or large-scale VPN.



On a Palo Alto Networks firewall or Panorama, you can import self-signed certificates only if they are CA certificates.

Instead of importing a self-signed root CA certificate into all the client systems, it is a best practice to import a certificate from the enterprise CA because the clients will already have a trust relationship with the enterprise CA, which simplifies the deployment.

If the certificate you will import is part of a certificate chain, it is a best practice to import the entire chain.

STEP 1 | From the enterprise CA, export the certificate and private key that the firewall will use for authentication.

When exporting a private key, you must enter a passphrase to encrypt the key for transport. Ensure the management system can access the certificate and key files. When importing the key onto the firewall, you must enter the same passphrase to decrypt it.

STEP 2 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 3 | If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

STEP 4 | Click **Import** and enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.

STEP 5 | To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the firewall supports multiple virtual systems.

© 2018 Palo Alto Networks, Inc.

STEP 6 | Enter the path and name of the **Certificate File** received from the CA, or **Browse** to find the file.

STEP 7 | Select a **File Format**:

- **Encrypted Private Key and Certificate (PKCS12)**—This is the default and most common format, in which the key and certificate are in a single container (**Certificate File**). If a hardware security module (HSM) will store the private key for this certificate, select the **Private key resides on Hardware Security Module** check box.
- **Base64 Encoded Certificate (PEM)**—You must import the key separately from the certificate. If a hardware security module (HSM) stores the private key for this certificate, select the **Private key resides on Hardware Security Module** check box and skip the next step. Otherwise, select the **Import Private Key** check box, enter the **Key File** or **Browse** to it, then continue to the next step.

STEP 8 | Enter and re-enter (confirm) the **Passphrase** used to encrypt the private key.

STEP 9 | Click **OK**. The Device Certificates page displays the imported certificate.

Obtain a Certificate from an External CA

The advantage of obtaining a certificate from an external certificate authority (CA) is that the private key does not leave the firewall. To obtain a certificate from an external CA, generate a certificate signing request (CSR) and submit it to the CA. After the CA issues a certificate with the specified attributes, import it onto the firewall. The CA can be a well-known, public CA or an enterprise CA.

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of the certificate, Configure an OCSP Responder before generating the CSR.

STEP 1 | Request the certificate from an external CA.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
2. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.
3. Click **Generate**.
4. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.
5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.
6. If the firewall has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box.
7. In the **Signed By** field, select **External Authority (CSR)**.
8. If applicable, select an **OCSP Responder**.
9. (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.

   If you add a Host Name attribute, it is a best practice for it to match the Common Name (this is mandatory for GlobalProtect). The host name populates the Subject Alternative Name field of the certificate.
10. Click **Generate**. The **Device Certificates** tab displays the CSR with a Status of `pending`.

STEP 2 | Submit the CSR to the CA.

1. Select the CSR and click **Export** to save the .csr file to a local computer.
2. Upload the .csr file to the CA.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Import the certificate.

1. After the CA sends a signed certificate in response to the CSR, return to the **Device Certificates** tab and click **Import**.
2. Enter the **Certificate Name** used to generate the CSR.
3. Enter the path and name of the PEM **Certificate File** that the CA sent, or **Browse** to it.
4. Click **OK**. The **Device Certificates** tab displays the certificate with a Status of `valid`.

STEP 4 | Configure the certificate.

1. Click the certificate **Name**.
2. Select the check boxes that correspond to the intended use of the certificate on the firewall. For example, if the firewall will use this certificate to secure forwarding of syslogs to an external syslog server, select the **Certificate for Secure Syslog** check box.
3. Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Export a Certificate and Private Key

Palo Alto Networks recommends that you use your enterprise public key infrastructure (PKI) to distribute a certificate and private key in your organization. However, if necessary, you can also export a certificate and private key from the firewall or Panorama. You can use an exported certificate and private key in the following cases:

- Configure Certificate-Based Administrator Authentication to the Web Interface
- Enable SSL Between GlobalProtect LSVPN Components to configure GlobalProtect agent/app authentication to portals and gateways
- SSL Forward Proxy decryption
- Obtain a Certificate from an External CA

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall has more than one virtual system (vsys), select a **Location** (a specific vsys or **Shared**) for the certificate.

STEP 3 | Select the certificate, click **Export**, and select a **File Format**:

- **Base64 Encoded Certificate (PEM)**—This is the default format. It is the most common and has the broadest support on the Internet. If you want the exported file to include the private key, select the **Export Private Key** check box.
- **Encrypted Private Key and Certificate (PKCS12)**—This format is more secure than PEM but is not as common or as broadly supported. The exported file will automatically include the private key.
- **Binary Encoded Certificate (DER)**—More operating system types support this format than the others. You can export only the certificate, not the key: ignore the **Export Private Key** check box and passphrase fields.

STEP 4 | Enter a **Passphrase** and **Confirm Passphrase** to encrypt the private key if the **File Format** is PKCS12 or if it is PEM and you selected the **Export Private Key** check box. You will use this passphrase when importing the certificate and key into client systems.

STEP 5 | Click **OK** and save the certificate/key file to your computer.

© 2018 Palo Alto Networks, Inc.

Configure a Certificate Profile

Certificate profiles define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPSec VPN, Mobile Security Manager, and web interface access to Palo Alto Networks firewalls or Panorama. The profiles specify which certificates to use, how to verify certificate revocation status, and how that status constrains access. Configure a certificate profile for each application.



It is a best practice to enable Online Certificate Status Protocol (OCSP) and/or Certificate Revocation List (CRL) status verification for certificate profiles. For details on these methods, see Certificate Revocation.

STEP 1 | Obtain the certificate authority (CA) certificates you will assign.

Perform one of the following steps to obtain the CA certificates you will assign to the profile. You must assign at least one.

- Generate a Certificate.
- Export a certificate from your enterprise CA and then import it onto the firewall (see Step Assign one or more certificates.).

STEP 2 | Identify the certificate profile.

1. Select **Device** > **Certificate Management** > **Certificates Profile** and click **Add**.
2. Enter a **Name** to identify the profile. The name is case-sensitive, must be unique and can use up to 31 characters that include only letters, numbers, spaces, hyphens, and underscores.
3. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

STEP 3 | Assign one or more certificates.

Perform the following steps for each CA certificate:

1. In the CA Certificates table, click **Add**.
2. Select a **CA Certificate**. Alternatively, to import a certificate, click **Import**, enter a **Certificate Name**, **Browse** to the **Certificate File** you exported from your enterprise CA, and click **OK**.
3. (**Optional**) If the firewall uses OCSP to verify certificate revocation status, configure the following fields to override the default behavior. For most deployments, these fields do not apply.
   - By default, the firewall uses the "Authority Information Access" (AIA) information from the certificate to extract the OCSP responder information. To override the AIA information, enter a **Default OCSP URL** (starting with **`http://`** or **`https://`**).
   - By default, the firewall uses the certificate selected in the **CA Certificate** field to validate OCSP responses. To use a different certificate for validation, select it in the **OCSP Verify CA Certificate** field.
4. Click **OK**. The CA Certificates table displays the assigned certificate.

STEP 4 | Define the methods for verifying certificate revocation status and the associated blocking behavior.

1. Select **Use CRL** and/or **Use OCSP**. If you select both, the firewall first tries OCSP and falls back to the CRL method only if the OCSP responder is unavailable.
2. Depending on the verification method, enter the **CRL Receive Timeout** and/or **OCSP Receive Timeout**. These are the intervals (1-60 seconds) after which the firewall stops waiting for a response from the CRL/OCSP service.

© 2018 Palo Alto Networks, Inc.

3. Enter the **Certificate Status Timeout**. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies any session-blocking logic you define. The **Certificate Status Timeout** relates to the OCSP/CRL **Receive Timeout** as follows:
   - If you enable both OCSP and CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the aggregate of the two **Receive Timeout** values.
   - If you enable only OCSP—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the OCSP **Receive Timeout** value.
   - If you enable only CRL—The firewall registers a request timeout after the lesser of two intervals passes: the **Certificate Status Timeout** value or the CRL **Receive Timeout** value.
4. If you want the firewall to block sessions when the OCSP or CRL service returns a certificate revocation status of unknown, select **Block session if certificate status is unknown**. Otherwise, the firewall allows the sessions.
5. If you want the firewall to block sessions after it registers an OCSP or CRL request timeout, select **Block session if certificate status cannot be retrieved within timeout**. Otherwise, the firewall allows the sessions.
6. (**GlobalProtect only**) If you want the firewall to block sessions when the serial number attribute in the subject of the client certificate does not match the host ID that the GlobalProtect agent reports for the client endpoint, select **Block sessions if the certificate was not issued to the authenticating device**.

STEP 5 | Click **OK** and **Commit**

© 2018 Palo Alto Networks, Inc.

Configure an SSL/TLS Service Profile

Palo Alto Networks firewalls and Panorama use 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 Captive Portal, GlobalProtect portals and gateways, inbound traffic on the management (MGT) interface, the URL Admin Override feature, and the User-ID™ syslog listening service. By defining the protocol versions, you can use a profile to restrict the cipher suites that are available for securing communication with the clients requesting the services. This improves network security by enabling the firewall or Panorama to avoid 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.

In the client systems that request firewall services, the certificate trust list (CTL) must include the certificate authority (CA) certificate that issued the certificate specified in the SSL/TLS service profile. Otherwise, users will see a certificate error when requesting firewall services. Most third-party CA certificates are present by default in client browsers. If an enterprise or firewall-generated CA certificate is the issuer, you must deploy that CA certificate to the CTL in client browsers.

STEP 1 | For each desired service, generate or import a certificate on the firewall (see Obtain Certificates).

Use only signed certificates, not CA certificates, in SSL/TLS service profiles.

STEP 2 | Select **Device** > **Certificate Management** > **SSL/TLS Service Profile**.

STEP 3 | If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available.

STEP 4 | Click **Add** and enter a **Name** to identify the profile.

STEP 5 | Select the **Certificate** you just obtained.

STEP 6 | Define the range of protocols that the service can use:

- For the **Min Version**, select the earliest allowed TLS version: **TLSv1.0** (default), **TLSv1.1**, or **TLSv1.2**.
- For the **Max Version**, select the latest allowed TLS version: **TLSv1.0**, **TLSv1.1**, **TLSv1.2**, or **Max** (latest available version). The default is **Max**.

As a best practice, set the Min Version to TLSv1.2 and the Max Version to Max.

On firewalls in FIPS/CC mode running PAN-OS 8.0 or a later release, TLSv1.1 is the earliest supported TLS version; do not select TLSv1.0.

Client certificates that are used when requesting firewall services that rely on TLSv1.2 cannot have SHA512 as a digest algorithm. The client certificates must use a lower digest algorithm (such as SHA384) or you must limit the Max Version to TLSv1.1 for the firewall services.

STEP 7 | Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Replace the Certificate for Inbound Management Traffic

When you first boot up the firewall or Panorama, it automatically generates a default certificate that enables HTTPS access to the web interface and XML API over the management (MGT) interface and (on the firewall only) over any other interface that supports HTTPS management traffic (for details, see Use Interface Management Profiles to Restrict Access). To improve the security of inbound management traffic, replace the default certificate with a new certificate issued specifically for your organization.



You cannot view, modify, or delete the default certificate.

To secure management traffic, you must also Configure Administrative Accounts and Authentication*.*

STEP 1 | Obtain the certificate that will authenticate the firewall or Panorama to the client systems of administrators.

You can simplify your Certificate Deployment by using a certificate that the client systems already trust. Therefore, we recommend that you Import a Certificate and Private Key from your enterprise certificate authority (CA) or Obtain a Certificate from an External CA; the trusted root certificate store of the client systems is likely to already have the associated root CA certificate that ensures trust.



If you Generate a Certificate *on the firewall or Panorama, administrators will see a certificate error because the root CA certificate is not in the trusted root certificate store of client systems. To prevent this, deploy the self-signed root CA certificate to all client systems.*



Regardless of how you obtain the certificate, we recommend a Digest algorithm of sha256 or higher for enhanced security.

STEP 2 | Configure an SSL/TLS Service Profile.

Select the **Certificate** you just obtained.



For enhanced security, we recommend that you set the Min Version (earliest allowed TLS version) to TLSv1.2 for inbound management traffic. We also recommend that you use a different SSL/TLS Service Profile for each firewall or Panorama service instead of reusing this profile for all services.

STEP 3 | Apply the SSL/TLS Service Profile to inbound management traffic.

1. Select **Device** > **Setup** > **Management** and edit the General Settings.
2. Select the **SSL/TLS Service Profile** you just configured.
3. Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Configure the Key Size for SSL Forward Proxy Server Certificates

When responding to a client in an SSL Forward Proxy session, the firewall creates a copy of the certificate that the destination server presents and uses the copy to establish a connection with the client. By default, the firewall generates certificates with the same key size as the certificate that the destination server presented. However, you can change the key size for the firewall-generated certificate as follows:

STEP 1 | Select **Device** > **Setup** > **Session** and, in the Decryption Settings section, click **SSL Forward Proxy Settings**.

STEP 2 | Select a **Key Size**:

- **Defined by destination host**—The firewall determines the key size for the certificates it generates to establish SSL proxy sessions with clients based on the key size of the destination server certificate. If the destination server uses a 1024-bit RSA key, the firewall generates a certificate with that key size and an SHA-1 hashing algorithm. If the destination server uses a key size larger than 1,024 bits (for example, 2,048 bits or 4,096 bits), the firewall generates a certificate that uses a 2,048-bit RSA key and SHA-256 algorithm. This is the default setting.
- **1024-bit RSA**—The firewall generates certificates that use a 1,024-bit RSA key and SHA-1 hashing algorithm regardless of the key size of the destination server certificates. As of December 31, 2013, public certificate authorities (CAs) and popular browsers have limited support for X.509 certificates that use keys of fewer than 2,048 bits. In the future, depending on security settings, when presented with such keys the browser might warn the user or block the SSL/TLS session entirely.
- **2048-bit RSA**—The firewall generates certificates that use a 2,048-bit RSA key and SHA-256 hashing algorithm regardless of the key size of the destination server certificates. Public CAs and popular browsers support 2,048-bit keys, which provide better security than the 1,024-bit keys.



Changing the key size setting clears the current certificate cache.

STEP 3 | Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Revoke and Renew Certificates

- Revoke a Certificate
- Renew a Certificate

Revoke a Certificate

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority (CA) that issued the certificate must revoke it. The following task describes how to revoke a certificate for which the firewall is the CA.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall supports multiple virtual systems, the tab displays a **Location** drop-down. Select the virtual system to which the certificate belongs.

STEP 3 | Select the certificate to revoke.

STEP 4 | Click **Revoke**. PAN-OS immediately sets the status of the certificate to revoked and adds the serial number to the Online Certificate Status Protocol (OCSP) responder cache or certificate revocation list (CRL). You need not perform a commit.

Renew a Certificate

If a certificate expires, or soon will, you can reset the validity period. If an external certificate authority (CA) signed the certificate and the firewall uses the Online Certificate Status Protocol (OCSP) to verify certificate revocation status, the firewall uses the OCSP responder information to update the certificate status (see Configure an OCSP Responder). If the firewall is the CA that issued the certificate, the firewall replaces it with a new certificate that has a different serial number but the same attributes as the old certificate.

STEP 1 | Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.

STEP 2 | If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.

STEP 3 | Select a certificate to renew and click **Renew**.

STEP 4 | Enter a **New Expiration Interval** (in days).

STEP 5 | Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Secure Keys with a Hardware Security Module

A hardware security module (HSM) is a physical device that manages digital keys. An HSM provides secure storage and generation of digital keys. It provides both logical and physical protection of these materials from non-authorized use and potential adversaries.

HSM clients integrated with Palo Alto Networks firewalls and Panorama enable enhanced security for the private keys used in SSL/TLS decryption (both SSL forward proxy and SSL inbound inspection). In addition, you can use the HSM to encrypt master keys.

The following topics describe how to integrate an HSM with your firewall or Panorama:

- Set up Connectivity with an HSM
- Encrypt a Master Key Using an HSM
- Store Private Keys on an HSM
- Manage the HSM Deployment

Set up Connectivity with an HSM

HSM clients are integrated with PA-3000 Series, PA-5000 Series, PA-5200 Series, PA-7000 Series, and VM-Series firewalls and with the Panorama management server (virtual appliance and M-Series appliances) for use with the following HSM vendors.

- **SafeNet Network**—The supported client versions depend on the PAN-OS release:
  - PAN-OS 8.0 releases and earlier releases—SafeNet Network client version 5.2.1.
  - PAN-OS 8.0.2 and later releases (also PAN-OS 7.1.10 and later PAN-OS 7.1 releases)—SafeNet Network client version 5.2.1, 5.4.2, and 6.2.2. On the firewall or Panorama, use the **`request hsm client-version`** CLI command to select the client version that is compatible with your SafeNet HSM server.
- **Thales nShield Connect**—All PAN-OS releases support client version 11.62.

The HSM server version must be compatible with these client versions. Refer to the HSM vendor documentation for the client-server version compatibility matrix.

- Set Up Connectivity with a SafeNet Network HSM
- Set Up Connectivity with a Thales nShield Connect HSM

Set Up Connectivity with a SafeNet Network HSM

To set up connectivity between the Palo Alto Networks firewall (HSM client) and a SafeNet Network HSM server, you must specify the IP address of the server, enter a password for authenticating the firewall to the server, and register the firewall with the server. Before starting the configuration, make sure you created a partition for the firewall on the HSM server. To ensure the SafeNet Network client version on the firewall is compatible with your SafeNet Network server, see Set up Connectivity with an HSM.

Before the HSM and firewall connect, the HSM authenticates the firewall based on the firewall IP address. Therefore, you must configure the firewall to use a static IP address, not a dynamic address assigned through DHCP. Operations on the HSM would stop working if the firewall IP address changed during runtime.



HSM configurations are not synchronized between high availability (HA) firewall peers. Consequently, you must configure the HSM separately on each peer. In active/passive HA deployments, you must manually perform one failover *to individually configure and*

© 2018 Palo Alto Networks, Inc.

authenticate each HA peer to the HSM. After this initial manual failover, user interaction is not required for the failover function.

STEP 1 | Define connection settings for each SafeNet Network HSM.

1. Log in to the firewall web interface and select **Device** > **Setup** > **HSM**.
2. Edit the Hardware Security Module Provider section and set the **Provider Configured** to **SafeNet Network HSM**.
3. **Add** each HSM server as follows. A high availability (HA) HSM configuration requires two servers.
   1. Enter a **Module Name** for the HSM server. This can be any ASCII string of up to 31 characters.
   2. Enter an IPv4 address for the HSM **Server Address**.
4. (**HA only**) Select **High Availability**, specify the **Auto Recovery Retry** value, and enter a **High Availability Group Name**.

   If two HSM servers are configured, the best practice is to enable High Availability. Otherwise the second HSM server is not used.
5. Click **OK** and **Commit**.

STEP 2 | (**Optional**) Configure a service route to connect to the HSM if you don't want the firewall to connect through the Management interface (default).

If you configure a service route for the HSM, running the ***`clear session all`*** *CLI command clears all existing HSM sessions, bringing all HSM states down and then up again. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.*

1. Select **Device** > **Setup** > **Services** and click **Service Route Configuration**.
2. **Customize** a service route. The **IPv4** tab is active by default.
3. Click **HSM** in the **Service** column.
4. Select a **Source Interface** for the HSM.
5. Click **OK** and **Commit**.

STEP 3 | Configure the firewall to authenticate to the HSM.

1. Select **Device** > **Setup** > **HSM** and **Setup Hardware Security Module**.
2. Select the HSM **Server Name**.
3. Enter the **Administrator Password** to authenticate the firewall to the HSM.
4. Click **OK**.

   The firewall tries to authenticate to the HSM and displays a status message.
5. Click **OK**.

STEP 4 | Register the firewall as an HSM client with the HSM server and assign the firewall to a partition on the HSM server.

If the HSM already has a firewall with the same <cl-name> registered, you must first remove the duplicate registration by running the ***`client delete -client <cl-name>`*** *command, where <cl-name> is the name of the client (firewall) registration you want to delete.*

1. Log in to the HSM from a remote system.
2. Register the firewall using the **`client register -c <cl-name> -ip <fw-ip-addr>`** command, where *<cl-name>* is a name that you assign to the firewall for use on the HSM and *<fw-ip-addr>* is the firewall IP address.

© 2018 Palo Alto Networks, Inc.

3. Assign a partition to the firewall using the **`client assignpartition -c <cl-name> -p <partition-name>`** command, where *<cl-name>* is the name assigned to the firewall in the **`client register`** command and *<partition-name>* is the name of a previously configured partition that you want to assign to the firewall.

STEP 5 | Configure the firewall to connect to the HSM partition.

1. Select **Device** > **Setup** > **HSM** and click the Refresh icon.
2. **Setup HSM Partition** in the Hardware Security Operations section.
3. Enter the **Partition Password** to authenticate the firewall to the partition on the HSM.
4. Click **OK**.

STEP 6 | (**HA only**) Repeat the previous authentication, registration, and partition connection steps to add another HSM to the existing HA group.

If you remove an HSM from your configuration, repeat the previous partition connection step to remove the deleted HSM from the HA group.

STEP 7 | Verify firewall connectivity and authentication with the HSM.

1. Select **Device** > **Setup** > **HSM** and check the authentication and connection Status:
   - **Green**—The firewall is successfully authenticated and connected to the HSM.
   - **Red**—The firewall failed to authenticate to the HSM or network connectivity to the HSM is down.
2. View the following columns in the Hardware Security Module Status section to determine the authentication status:
   - **Serial Number**—The serial number of the HSM partition if the firewall successfully authenticated to the HSM.
   - **Partition**—The partition name on the HSM that is assigned on the firewall.
   - **Module State**—The current state of the HSM connection. The value is always `Authenticated` if the Hardware Security Module Status section displays the HSM.

Set Up Connectivity with a Thales nShield Connect HSM

You must set up a remote filesystem (RFS) as a hub to synchronize key data for all the firewalls (HSM clients) in your organization that use the Thales nShield Connect HSM. To ensure the Thales nShield Connect client version on your firewalls is compatible with your Thales nShield Connect server, see Set up Connectivity with an HSM.

Before the HSM and firewalls connect, the HSM authenticates the firewalls based on their IP addresses. Therefore, you must configure the firewalls to use static IP addresses, not dynamic addresses assigned through DHCP. Operations on the HSM would stop working if the firewall IP addresses changed during runtime.

HSM configurations are not synchronized between high availability (HA) firewall peers. Consequently, you must configure the HSM separately on each peer. In active/passive HA deployments, you must manually perform one failover to individually configure and authenticate each HA peer to the HSM. After this initial manual failover, user interaction is not required for the failover function.

STEP 1 | Define connection settings for each Thales nShield Connect HSM.

1. Log in to the firewall web interface and select **Device** > **Setup** > **HSM**.

© 2018 Palo Alto Networks, Inc.

2. Edit the Hardware Security Module Provider section and set the **Provider Configured** to **Thales nShield Connect**.
3. **Add** each HSM server as follows. A high availability (HA) HSM configuration requires two servers.
    1. Enter a **Module Name** for the HSM server. This can be any ASCII string of up to 31 characters.
    2. Enter an IPv4 address for the HSM **Server Address**.
4. Enter an IPv4 address for the **Remote Filesystem Address**.
5. Click **OK** and **Commit**.

STEP 2 | (**Optional**) Configure a service route to connect to the HSM if you don't want the firewall to connect through the Management interface (default).



If you configure a service route for the HSM, running the ***clear session all*** *CLI command clears all existing HSM sessions, bringing all HSM states down and then up again. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.*

1. Select **Device** > **Setup** > **Services** and click **Service Route Configuration**.
2. **Customize** a service route. The **IPv4** tab is active by default.
3. Click **HSM** in the **Service** column.
4. Select a **Source Interface** for the HSM.
5. Click **OK** and **Commit**.

STEP 3 | Register the firewall as an HSM client with the HSM server.

This step briefly describes the procedure for using the front panel interface of the Thales nShield Connect HSM. For more details, refer to the Thales documentation.

1. Log in to the front panel display of the Thales nShield Connect HSM unit.
2. Use the right-hand navigation button to select **System** > **System configuration** > **Client config** > **New client**.
3. Enter the firewall IP address.
4. Select **System** > **System configuration** > **Client config** > **Remote file system** and enter the IP address of the client computer where you set up the RFS.

STEP 4 | Configure the RFS to accept connections from the firewall.

1. Log in to the RFS from a Linux client.
2. Obtain the electronic serial number (ESN) and the hash of the K_{NETI} key, which authenticates the HSM to clients, by running the **anonkneti *<ip-address>*** command, where *<ip-address>* is the HSM IP address.

    The following is an example:

    ```
    anonkneti 192.0.2.1
    ```

    ```
    B1E2-2D4C-E6A2 5a2e5107e70d525615a903f6391ad72b1c03352c
    ```

    In this example, `B1E2-2D4C-E6A2` is the ESN and `5a2e5107e70d525615a903f6391ad72b1c03352c` is the hash of the K_{NETI} key.
3. Use the following command from a superuser account to set up the RFS:

    ```
    rfs-setup --force <ip-address> <ESN> <hash-Kneti-key>
    ```


© 2018 Palo Alto Networks, Inc.

The *<ip-address>* is the IP address of the HSM, *<ESN>* is the electronic serial number, and *<hash-Kneti-key>* is the hash of the K_{NETI} key.

The following example uses the values obtained in this procedure:

```
rfs-setup --force 192.0.2.1 B1E2-2D4C-E6A2
 5a2e5107e70d525615a903f6391ad72b1c03352c
```

4. Use the following command to permit HSM client submissions on the RFS:

```
rfs-setup --gang-client --write-noauth <FW-IPaddress>
```

where *<FW-IPaddress>* is the firewall IP address.

STEP 5 | Authenticate the firewall to the HSM.

1. In the firewall web interface, select **Device** > **Setup** > **HSM** and **Setup Hardware Security Module**.
2. Click **OK**.

   The firewall tries to authenticate to the HSM and displays a status message.
3. Click **OK**.

STEP 6 | Synchronize the firewall with the RFS by selecting **Device** > **Setup** > **HSM** and **Synchronize with Remote Filesystem**.

STEP 7 | Verify firewall connectivity and authentication with the HSM.

1. Select **Device** > **Setup** > **HSM** and check the authentication and connection Status:
   - **Green**—The firewall is successfully authenticated and connected to the HSM.
   - **Red**—The firewall failed to authenticate to the HSM or network connectivity to the HSM is down.
2. Check the Hardware Security Module Status section to determine the authentication status.
   - **Name**—The name of the HSM.
   - **IP address**—The IP address of the HSM.
   - **Module State**—The current state of the HSM connection: `Authenticated` or `Not Authenticated`.

Encrypt a Master Key Using an HSM

A master key encrypts all private keys and passwords on the firewall and Panorama. If you have security requirements to store your private keys in a secure location, you can encrypt the master key using an encryption key that is stored on an HSM. The firewall or Panorama then requests the HSM to decrypt the master key whenever it is required to decrypt a password or private key on the firewall. Typically, the HSM is in a highly secure location that is separate from the firewall or Panorama for greater security.

The HSM encrypts the master key using a wrapping key. To maintain security, you must occasionally change (refresh) this wrapping key.

Firewalls configured in FIPS/CC mode do not support master key encryption using an HSM.

The following topics describe how to encrypt the master key initially and how to refresh the master key encryption:

- Encrypt the Master Key
- Refresh the Master Key Encryption

© 2018 Palo Alto Networks, Inc.

Encrypt the Master Key

If you have not previously encrypted the master key on a firewall, use the following procedure to encrypt it. Use this procedure for first time encryption of a key, or if you define a new master key and you want to encrypt it. If you want to refresh the encryption on a previously encrypted key, see Refresh the Master Key Encryption.

STEP 1 | Select **Device** > **Master Key and Diagnostics**.

STEP 2 | Specify the key that is currently used to encrypt all of the private keys and passwords on the firewall in the **Master Key** field.

STEP 3 | If changing the master key, enter the new master key and confirm.

STEP 4 | Select the **HSM** check box.

- **Life Time**—The number of days and hours after which the master key expires (range 1-730 days).
- **Time for Reminder**—The number of days and hours before expiration when the user is notified of the impending expiration (range 1–365 days).

STEP 5 | Click **OK**.

Refresh the Master Key Encryption

As a best practice, periodically refresh the master key encryption by rotating the wrapping key that encrypts it. The frequency of the rotation depends on your application. The wrapping key resides on your HSM. The following command is the same for SafeNet Network and Thales nShield Connect HSMs.

STEP 1 | Log in to the firewall CLI.

STEP 2 | Use the following CLI command to rotate the wrapping key for the master key on an HSM:

```
> request hsm mkey-wrapping-key-rotation
```

If the master key is encrypted on the HSM, the CLI command will generate a new wrapping key on the HSM and encrypt the master key with the new wrapping key.

If the master key is not encrypted on the HSM, the CLI command will generate new wrapping key on the HSM for future use.

The old wrapping key is not deleted by this command.

Store Private Keys on an HSM

For added security, you can use an HSM to secure the private keys used in SSL/TLS decryption for:

- SSL Forward Proxy—The HSM can store the private key of the Forward Trust certificate that signs certificates in SSL/TLS forward proxy operations. The firewall will then send the certificates that it generates during such operations to the HSM for signing before forwarding the certificates to the client.
- SSL Inbound Inspection—The HSM can store the private keys for the internal servers for which you are performing SSL/TLS inbound inspection.

If you use the DHE or ECDHE key exchange algorithms to enable Perfect Forward Secrecy (PFS) Support for SSL Decryption, you cannot use an HSM to store the private keys for SSL

© 2018 Palo Alto Networks, Inc.

Inbound Inspection. You also cannot use an HSM to store ECDSA keys used for Forward Proxy or Inbound Inspection decryption.

STEP 1 | On the HSM, import or generate the certificate and private key used in your decryption deployment.

For instructions on importing or generating a certificate and private key on the HSM, refer to your HSM documentation.

STEP 2 | (Thales nShield Connect only) Synchronize the key data from the Thales nShield remote file system to the firewall.



Synchronization with the SafeNet Network HSM is automatic.

1. Access the firewall web interface and select **Device** > **Setup** > **HSM**.
2. Select **Synchronize with Remote Filesystem** in the Hardware Security Operations section.

STEP 3 | Import the certificate that corresponds to the HSM-stored key onto the firewall.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
2. Enter the **Certificate Name**.
3. **Browse** to the **Certificate File** on the HSM.
4. Select a **File Format**.
5. Select **Private Key resides on Hardware Security Module**.
6. Click **OK** and **Commit**.

STEP 4 | (Forward Trust certificates only) Enable the certificate for use in SSL/TLS Forward Proxy.

1. Open the certificate you imported in Step 3 for editing.
2. Select **Forward Trust Certificate**.
3. Click **OK** and **Commit**.

STEP 5 | Verify that you successfully imported the certificate onto the firewall.

Locate the certificate you imported in Step 3 and check the icon in the Key column:

- **Lock icon**—The private key for the certificate is on the HSM.
- **Error icon**—The private key is not on the HSM or the HSM is not properly authenticated or connected.

Manage the HSM Deployment

You can perform the following tasks to manage your HSM deployment:

- View the HSM configuration settings.

  Select **Device** > **Setup** > **HSM**.

- Display detailed HSM information.

  Select **Show Detailed Information** from the Hardware Security Operations section.

  Information regarding the HSM servers, HSM HA status, and HSM hardware is displayed.

- Export Support file.

© 2018 Palo Alto Networks, Inc.

Select **Export Support File** from the Hardware Security Operations section.

A test file is created to help customer support when addressing a problem with an HSM configuration on the firewall.

- Reset HSM configuration.

  Select **Reset HSM Configuration** from the Hardware Security Operations section.

  Selecting this option removes all HSM connections. All authentication procedures must be repeated after using this option.

High Availability

High availability (HA) is a deployment in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point of failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up two firewalls in an HA pair provides redundancy and allows you to ensure business continuity.

> HA Overview
> HA Concepts
> Set Up Active/Passive HA
> Set Up Active/Active HA
> HA Firewall States
> Reference: HA Synchronization

© 2018 Palo Alto Networks, Inc.

HA Overview

You can set up two Palo Alto Networks firewalls as an HA pair. HA allows you to minimize downtime by making sure that an alternate firewall is available in the event that the peer firewall fails. The firewalls in an HA pair use dedicated or in-band HA ports on the firewall to synchronize data—network, object, and policy configurations—and to maintain state information. Firewall-specific configuration such as management interface IP address or administrator profiles, HA specific configuration, log data, and the Application Command Center (ACC) information is not shared between peers. For a consolidated application and log view across the HA pair, you must use Panorama, the Palo Alto Networks centralized management system.

When a failure occurs on a firewall in an HA pair and the 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. (Link Monitoring)
- One or more of the destinations specified on the firewall cannot be reached. (Path Monitoring)
- The firewall does not respond to heartbeat polls. (Heartbeat Polling and Hello messages)
- A critical chip or software component fails, known as packet path health monitoring.

You can use Panorama to manage HA firewalls. See Context Switch—Firewall or Panorama in the Panorama Administrator's Guide.

Palo Alto Networks firewalls support stateful active/passive or active/active high availability with session and configuration synchronization with a few exceptions:

- The PA-200 firewall supports HA Lite only. HA Lite is an active/passive deployment that provides configuration synchronization and some runtime data synchronization such as IPSec security associations. It does not support any session synchronization (HA2), and therefore does not offer stateful failover.
- The VM-Series firewall in AWS supports active/passive HA only; if it 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 does not support HA.

After you understand the HA Concepts, proceed to Set Up Active/Passive HA or Set Up Active/Active HA.

© 2018 Palo Alto Networks, Inc.

HA Concepts

The following topics provide conceptual information about how HA works on a Palo Alto Networks firewall:

- HA Modes
- HA Links and Backup Links
- Device Priority and Preemption
- Failover
- LACP and LLDP Pre-Negotiation for Active/Passive HA
- Floating IP Address and Virtual MAC Address
- ARP Load-Sharing
- Route-Based Redundancy
- HA Timers
- Session Owner
- Session Setup
- NAT in Active/Active HA Mode
- ECMP in Active/Active HA Mode

HA Modes

You can set up the firewalls for HA 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.

  The PA-200 firewall supports HA Lite only. HA Lite is an active/passive deployment that provides configuration synchronization and some runtime data synchronization such as IPSec security associations. It does not support any session synchronization (HA2), and therefore does not offer stateful failover.

- **Active/Active**— Both firewalls in the pair are active and processing traffic and work synchronously to handle session setup and session ownership. Both firewalls individually maintain session tables and routing tables and synchronize to each other. Active/active HA is supported in virtual wire and Layer 3 deployments.

  In active/active HA mode, the firewall does not support 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.

  An active/active configuration does not load-balance traffic. Although you can load-share by sending traffic to the peer, no load balancing occurs. Ways to load share sessions to both firewalls include using ECMP, multiple ISPs, and load balancers.

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 routing and traffic flow issues in active/passive mode. Active/passive mode supports a Layer 2 deployment; 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

© 2018 Palo Alto Networks, Inc.

addresses to provide proper failover. Because both firewalls are actively processing traffic, the firewalls use additional concepts of session owner and session setup to perform Layer 7 content inspection. 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 are actively processing traffic.



In active/active mode, the HA pair can be used to temporarily process more traffic than what one firewall can normally handle. However, this should not be the norm because a failure of one firewall causes all traffic to be redirected to the remaining firewall in the HA pair. Your design must allow the remaining firewall to process the maximum capacity of your traffic loads with content inspection enabled. If the design oversubscribes the capacity of the remaining firewall, high latency and/or application failure can occur.

For information on setting up your firewalls in active/passive mode, see Set Up Active/Passive HA. For information on setting up your firewalls in active/active mode, see Set Up Active/Active HA.

HA Links and Backup 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 details, see HA Ports on Palo Alto Networks Firewalls.
- For firewalls without dedicated HA ports such as the PA-200, PA-220, and PA-500 firewalls, as a best practice use the dataplane port for the HA port, and use the management port as the HA1 backup.

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).
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.
Backup Links	Provide redundancy for the HA1 and the HA2 links. In-band ports are used as backup links for both HA1 and HA2. Consider the following guidelines when configuring backup HA links:

© 2018 Palo Alto Networks, Inc.

HA Links and Backup Links	Description
	• 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. *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-3000 Series, PA-5000 Series, and PA-5200 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-5200 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.

HA Ports on Palo Alto Networks Firewalls

When connecting two Palo Alto Networks® firewalls in a high availability (HA) configuration, we recommend that you use the dedicated HA ports for HA Links and Backup Links. These dedicated ports include: the HA1 ports labeled HA1, HA1-A, and HA1-B used for HA control and synchronization traffic; and HA2 and the High Speed Chassis Interconnect (HSCI) ports used for HA session setup traffic. The PA-5200 Series firewalls have multipurpose auxiliary ports labeled AUX-1 and AUX-2 that you can configure for HA1 traffic.

You can also configure the HSCI port for HA3, which is used for packet forwarding to the peer firewall during session setup and asymmetric traffic flow (active/active HA only). The HSCI port can be used for HA2 traffic, HA3 traffic, or both.

The HA1, HA2, and AUX links provide synchronization for functions that reside on the management plane. Using the dedicated HA interfaces on the management plane is more efficient than using the in-band ports as this eliminates the need to pass the synchronization packets over the dataplane.

If your firewall does not have dedicated HA ports, you can configure data ports as HA interfaces. If your firewall does have dedicated HA ports but does not have a dedicated HA backup port, you can also configure data ports as backups to dedicated HA ports.

Whenever possible, connect HA ports directly between the two firewalls in an HA pair (not through a switch or router) to avoid HA link and communications problems that could occur if there is a network issue.

© 2018 Palo Alto Networks, Inc.

Use the following table to learn about dedicated HA ports and how to connect the HA Links and Backup Links:

Model	Front-Panel Dedicated Port(s)

PA-800 Series, PA-3000 Series, and PA-5000 Series Firewalls	• **HA1 and HA2**—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 and HA2 in both HA Modes.
	  • **For HA1 traffic**—Connect the HA1 port on the first firewall directly to the HA1 port on the second firewall in the pair or connect these ports together through a switch or router.
	  • **For HA2 traffic**—Connect the HA2 port on the first firewall directly to the HA2 port on the second firewall in the pair or connect these ports together through a switch or router.

PA-5200 Series Firewalls	• **HA1-A and HA1-B**—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 traffic in both HA Modes.
	  • **For HA1 traffic**—Connect the HA1-A port on the first firewall directly to the HA1-A port on the second firewall in the pair or connect them together through a switch or router.
	  • **For a backup to the HA1-A connection**—Connect the HA1-B port on the first firewall directly to the HA1-B port on the second firewall in the pair or connect them together through a switch or router.
	• **HSCI**—The HSCI port is a Layer 1 interface that connects two PA-5200 Series firewalls in an HA configuration. Use this port for an HA2 connection, HA3 connection, or both.
	
	The HSCI port on the PA-5220 firewall is a QSFP+ port and the HSCI port on the PA-5250, PA-5260, and PA-5280 firewalls is a QSFP28 port.
	
	The traffic carried on the HSCI port is raw Layer 1 traffic, which is not routable or switchable. Therefore, you must connect the HSCI ports directly to each other (from the HSCI port on the first firewall to the HSCI port on the second firewall).

PA-5200 Series Firewalls (continued)	• **AUX-1 and AUX-2**—The auxiliary SFP+ ports are multipurpose ports that you can configure for HA1, management functions, or log forwarding to Panorama. Use these ports when you need a fiber connection for one of these functions.
	  • **For HA1 traffic**—Connect the AUX-1 port on the first firewall directly to the AUX-1 port on the second firewall in the pair or connect them together through a switch or router.
	  • **For a backup to the AUX-1 connection**—Connect the AUX-2 port on the first firewall directly to the AUX-2 port on the second firewall in the pair or connect them together through a switch or router.

PA-7000 Series Firewalls *For HA2 or HA2/ HA3 traffic, the*	• **HA1-A and HA1-B**—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 traffic in both HA Modes. • **For HA1 traffic**—Connect the HA1-A port on the first firewall directly to the HA1-A port on the second firewall in the pair or connect them together through a switch or router.

© 2018 Palo Alto Networks, Inc.

Model	Front-Panel Dedicated Port(s)
PA-7000 Series firewalls synchronize sessions across the NPCs one-for-one.	• **For a backup to the HA1-A connection**—Connect the HA1-B port on the first firewall directly to the HA1-B port on the second firewall in the pair or connect them together through a switch or router. *You cannot configure an HA1 connection on the NPC data ports or the management (MGT) port.* • **HSCI-A and HSCI-B**—The HSCI ports are Layer 1 QSFP+ interfaces that connect two PA-7000 Series firewalls in an HA configuration. Use these ports for an HA2 connection, HA3 connection, or both. The traffic carried on the HSCI ports is raw Layer 1 traffic, which is not routable or switchable. Therefore, you must connect these ports as follows: • **For HA2 and HA3 traffic**—Connect the HSCI-A port on the first firewall directly to the HSCI-A port on the second firewall. • **For a backup to the HSCI-A connection**—Connect the HSCI-B port on the first firewall directly to the HSCI-B port on the second firewall.

Device Priority and Preemption

The firewalls in an HA pair can be assigned a *device priority* value to indicate a preference for which firewall should assume the active or active-primary role. If you need to use a specific firewall in the HA pair for actively securing traffic, you must enable the preemptive behavior on both the firewalls and assign a device priority value for each firewall. The firewall with the lower numerical value, and therefore *higher priority*, is designated as active or active-primary. The other firewall is the active-secondary or passive firewall.

By default, preemption is disabled on the firewalls and must be enabled on both firewalls. When enabled, the preemptive behavior allows the firewall with the *higher priority* (lower numerical value) to resume as active or active-primary after it recovers from a failure. When preemption occurs, the event is logged in the system logs.

Failover

When a failure occurs on one firewall and the peer 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 are monitored 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. By default, the interval for the heartbeat is 1000 milliseconds. A ping is sent every 1000 milliseconds and if there are three consecutive heartbeat losses, a failovers occurs. For details on the HA timers that trigger a failover, see HA Timers on page 260.
- **Link Monitoring**

  The physical interfaces to be monitored are grouped into a link group and their state (link up or link down) is monitored. A link group can contain one or more physical interfaces. A firewall failure is triggered when any or all of the interfaces in the group fail. The default behavior is failure of any one link

© 2018 Palo Alto Networks, Inc.

in the link group will cause the firewall to change the HA state to non-functional (or to tentative state in active/active mode) to indicate a failure of a monitored object.

- **Path Monitoring**

  Monitors the full path through the network to mission-critical IP addresses. ICMP pings are used to verify reachability of the IP address. The default interval for pings is 200ms. An IP address is considered unreachable when 10 consecutive pings (the default value) fail, and a firewall failure is triggered when any or all of the IP addresses monitored become unreachable. The default behavior is any one of the IP addresses becoming unreachable will cause the firewall to change the HA state to non-functional (or to tentative state in active/active mode) to indicate a failure of a monitored object.

In addition to the failover triggers listed above, a failover also occurs when the administrator suspends the firewall or when preemption occurs.

On the PA-3000 Series, PA-5000 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.

LACP and LLDP Pre-Negotiation for Active/Passive HA

If a firewall uses LACP or LLDP, negotiation of those protocols upon failover prevents sub-second failover. However, you can enable an interface on a passive firewall to negotiate LACP and LLDP prior to failover. Thus, a firewall in Passive or Non-functional HA state can communicate with neighboring devices using LACP or LLDP. Such pre-negotiation speeds up failover.

The PA-3000 Series, PA-5000 Series, PA-5200 Series, and PA-7000 Series firewalls support a pre-negotiation configuration depending on whether the Ethernet or AE interface is in a Layer 2, Layer 3, or virtual wire deployment. An HA passive firewall handles LACP and LLDP packets in one of two ways:

- **Active**—The firewall has LACP or LLDP configured on the interface and actively participates in LACP or LLDP pre-negotiation, respectively.
- **Passive**—LACP or LLDP is not configured on the interface and the firewall does not participate in the protocol, but allows the peers on either side of the firewall to pre-negotiate LACP or LLDP, respectively.

Pre-negotiation is not supported on subinterfaces or tunnel interfaces.

To configure LACP or LLDP pre-negotiation, see Step (Optional) Enable LACP and LLDP Pre-Negotiation for Active/Passive HA for faster failover if your network uses LACP or LLDP.

Floating IP Address and Virtual MAC Address

In a Layer 3 deployment of HA active/active mode, you can assign floating IP addresses, which 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 Virtual Router Redundancy Protocol (VRRP). Floating IP addresses can also be used to implement VPNs and source NAT, allowing for persistent connections when a firewall offering those services fails.

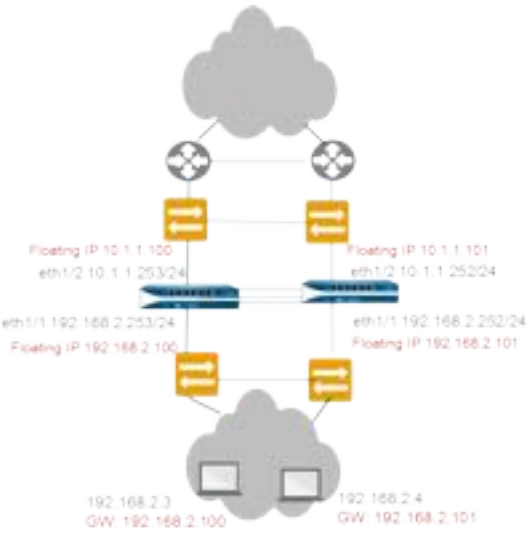
As shown in the figure below, each HA firewall interface has its own IP address and floating IP address. The interface IP address remains local to the firewall, but the floating IP address moves between the firewalls upon firewall failure. You configure the end hosts to use a floating IP address as its default gateway, allowing you to load balance 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 below, each firewall has two floating IP addresses and virtual MAC addresses; they all move over if the firewall fails.) The functioning firewall sends

© 2018 Palo Alto Networks, Inc.

a gratuitous ARP to update the MAC tables of the connected switches to inform them of the change in floating IP address and MAC address ownership to redirect traffic to itself.

After the failed firewall recovers, by default the floating IP address and virtual MAC address move back to firewall with the Device ID [0 or 1] to which the floating IP address is bound. More specifically, after the failed firewall recovers, it comes on line. The currently active firewall determines that the firewall is back online and checks whether the floating IP address it is handling belongs natively to itself or the other firewall. If the floating IP address was originally bound to the other Device ID, the firewall automatically gives it back. (For an alternative to this default behavior, see Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall.)



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.

The format of the virtual MAC address (on firewalls other than PA-7000 Series firewalls) is 00-1B-17-00-xx-yy, where 00-1B-17 is the vendor ID (of Palo Alto Networks in this case), 00 is fixed, xx indicates the Device ID and Group ID as shown in the following figure, and yy is the Interface ID:



The format of the virtual MAC address on PA-7000 Series firewalls is 00-1B-17-xx-xx-xx, where 00-1B-17 is the vendor ID (of Palo Alto Networks in this case), and the next 24 bits indicate the Device ID, Group ID and Interface ID as follows:

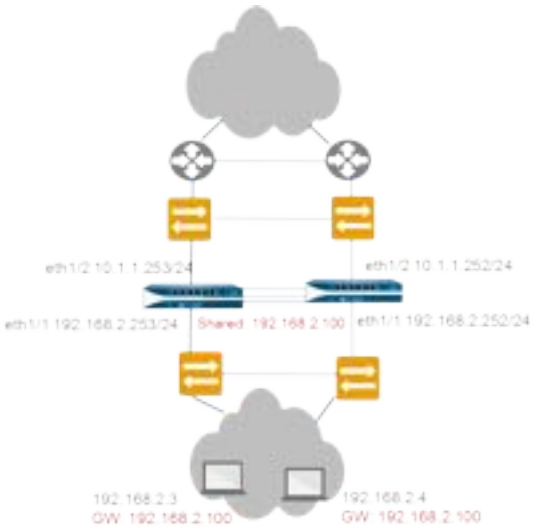


When a new active firewall takes over, it sends gratuitous ARPs from each of its connected interfaces to inform the connected Layer 2 switches of the new location of the virtual MAC address. To configure floating IP addresses, see Use Case: Configure Active/Active HA with Floating IP Addresses.

© 2018 Palo Alto Networks, Inc.

ARP Load-Sharing

In a Layer 3 interface deployment and active/active HA configuration, ARP load-sharing allows the firewalls to share an IP address and provide gateway services. Use ARP load-sharing only when no Layer 3 device exists between the firewall and end hosts, that is, when end hosts use the firewall as their default gateway.



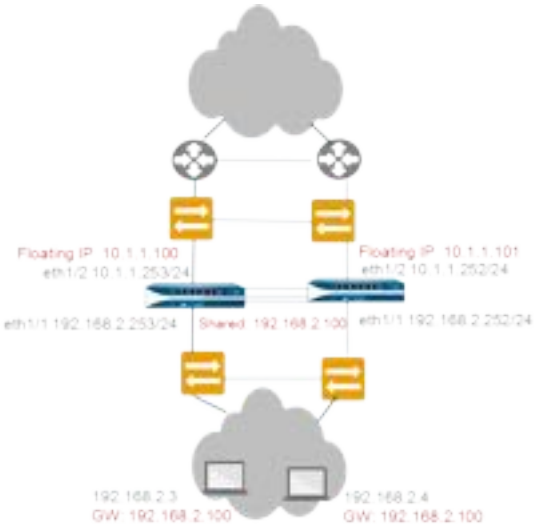
In such a scenario, all hosts are configured with a single gateway IP address. One of the firewalls responds to ARP requests for the gateway IP address with its virtual MAC address. Each firewall has a unique virtual MAC address generated for the shared IP address. The load-sharing algorithm that controls which firewall will respond to the ARP request is configurable; it is determined by computing the hash or modulo of the source IP address of the ARP request.

After the end host receives the ARP response from the gateway, it caches the MAC address and all traffic from the host is routed via the firewall that responded with the virtual MAC address for the lifetime of the ARP cache. The lifetime of the ARP cache depends on the end host operating system.

If a link or firewall fails, the floating IP address and virtual MAC address move over to the functional firewall. The functional firewall sends gratuitous ARPs to update the MAC table of the connected switches to redirect traffic from the failed firewall to itself. See Use Case: Configure Active/Active HA with ARP Load-Sharing.

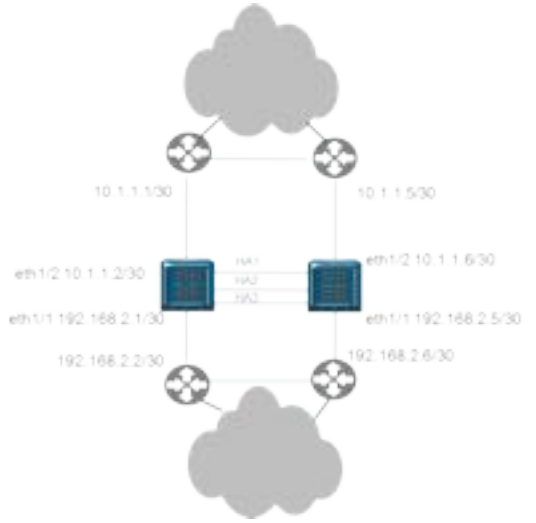
You can configure interfaces on the WAN side of the HA firewalls with floating IP addresses, and configure interfaces on the LAN side of the HA firewalls with a shared IP address for ARP load-sharing. For example, the figure below illustrates floating IP addresses for the upstream WAN edge routers and an ARP load-sharing address for the hosts on the LAN segment.

© 2018 Palo Alto Networks, Inc.



Route-Based Redundancy

In a Layer 3 interface deployment and active/active HA configuration, the firewalls are connected to routers, not switches. The firewalls use dynamic routing protocols to determine the best path (asymmetric route) and to load share between the HA pair. In such a scenario, no floating IP addresses are necessary. If a link, monitored path, or firewall fails, or if Bidirectional Forwarding Detection (BFD) detects a link failure, the routing protocol (RIP, OSPF, or BGP) handles the rerouting of traffic to the functioning firewall. You configure each firewall interface with a unique IP address. The IP addresses remain local to the firewall where they are configured; they do not move between devices when a firewall fails. See Use Case: Configure Active/Active HA with Route-Based Redundancy.



HA Timers

High availability (HA) timers facilitate a firewall to detect a firewall failure and trigger a failover. To reduce the complexity in configuring HA timers, you can select from three profiles: **Recommended**, **Aggressive** and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.

© 2018 Palo Alto Networks, Inc.

Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

The following table describes each timer included in the profiles and the current preset values (Recommended/Aggressive) across the different hardware models; these values are for current reference only and can change in a subsequent release.

Timers	Description	PA-7000 Series PA-5200 Series PA-5000 Series PA-3000 Series VM-Series	PA-800 Series PA-500 PA-220 PA-200	Panorama Virtual Appliance Panorama M-Series
Monitor Fail Hold Up Time (ms)	Interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is recommended to avoid an HA failover due to the occasional flapping of neighboring devices.	0/0	0/0	0/0
Preemption Hold Time (min)	Time that a passive or active-secondary firewall will wait before taking over as the active or active-primary firewall.	1/1	1/1	1/1
Heartbeat Interval (ms)	Frequency at which the HA peers exchange heartbeat messages in the form of an ICMP (ping).	1000/1000 2000/1000 (only for VM-Series in AWS)	2000/1000	2000/1000
Promotion Hold Time (ms)	Time that the passive firewall (in active/passive mode) or the active-secondary firewall (in active/active mode) will wait before taking over as the active or active-primary firewall after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made.	2000/500	2000/500	2000/500

© 2018 Palo Alto Networks, Inc.

Timers	Description	PA-7000 Series PA-5200 Series PA-5000 Series PA-3000 Series VM-Series	PA-800 Series PA-500 PA-220 PA-200	Panorama Virtual Appliance Panorama M-Series
Additional Master Hold Up Time (ms)	Time interval that is applied to the same event as Monitor Fail Hold Up Time (range 0-60000 ms, default 500 ms). The additional time interval is applied only to the active firewall in active/passive mode and to the active-primary firewall in active/active mode. This timer is recommended to avoid a failover when both firewalls experience the same link/path monitor failure simultaneously.	500/500	500/500	7000/5000
Hello Interval (ms)	Interval in milliseconds between hello packets that are sent to verify that the HA functionality on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms.	8000/8000	8000/8000	8000/8000
Maximum No. of Flaps	A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. This value indicates the maximum number of flaps that are permitted before the firewall is determined to be suspended and the passive firewall takes over (range 0-16; default 3).	3/3	3/3	Not Applicable

© 2018 Palo Alto Networks, Inc.

Session Owner

In an HA active/active configuration, both firewalls are active simultaneously, which means packets can be distributed between them. Such distribution requires the firewalls to fulfill two functions: session ownership and session setup. Typically, each firewall of the pair performs one of these functions, thereby avoiding race conditions that can occur in asymmetrically routed environments.

You configure the session owner of sessions to be either the firewall that receives the First Packet of a new session from the end host or the firewall that is in active-primary state (the Primary device). If Primary device is configured, but the firewall that receives the first packet is not in active-primary state, the firewall forwards the packet to the peer firewall (the session owner) over the HA3 link.

The session owner performs all Layer 7 processing, such as App-ID, Content-ID, and threat scanning for the session. The session owner also generates all traffic logs for the session.

If the session owner fails, the peer firewall becomes the session owner. The existing sessions fail over to the functioning firewall and no Layer 7 processing is available for those sessions. When a firewall recovers from a failure, by default, all sessions it owned before the failure revert back to that original firewall; Layer 7 processing does not resume.

If you configure session ownership to be Primary device, the session setup defaults to Primary device also.



Palo Alto Networks recommends setting the Session Owner to First Packet and the Session Setup to IP Modulo unless otherwise indicated in a specific use case.



Setting Session Owner and Session Setup to Primary Device causes the active-primary firewall to perform all traffic processing. You might want to configure this for one of these reasons:

- *You are troubleshooting and capturing logs and pcaps, so that packet processing is not split between the firewalls.*
- *You want to force the active/active HA pair to function like an active/passive HA pair. See* Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall *on page 285.*

Session Setup

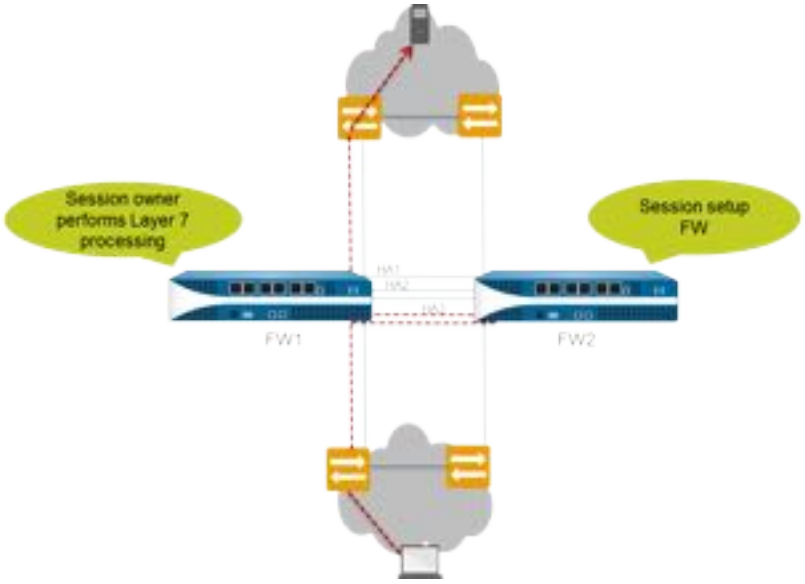
The session setup firewall performs the Layer 2 through Layer 4 processing necessary to set up a new session. The session setup firewall also performs NAT using the NAT pool of the session owner. You determine the session setup firewall in an active/active configuration by selecting one of the following session setup load sharing options.

Session Setup Option	Description
IP Modulo	The firewall distributes the session setup load based on parity of the source IP address. This is a deterministic method of sharing the session setup.
IP Hash	The firewall uses a hash of the source and destination IP addresses to distribute session setup responsibilities.
Primary Device	The active-primary firewall always sets up the session; only one firewall performs all session setup responsibilities.
First Packet	The firewall that receives the first packet of a session performs session setup.

© 2018 Palo Alto Networks, Inc.

- *If you want to load-share the session owner and session setup responsibilities, set session owner to First Packet and session setup to IP modulo. These are the recommended settings.*
- *If you want to do troubleshooting or capture logs or pcaps, or if you want an active/active HA pair to function like an active/passive HA pair, set both the session owner and session setup to Primary device so that the active-primary device performs all traffic processing. See* Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall*.*

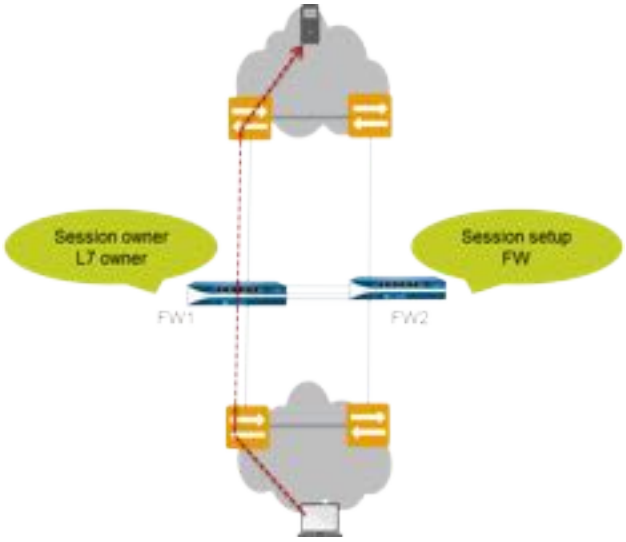
The firewall uses the HA3 link to send packets to its peer for session setup if necessary. The following figure and text describe the path of a packet that firewall FW1 receives for a new session. The red dotted lines indicate FW1 forwarding the packet to FW2 and FW2 forwarding the packet back to FW1 over the HA3 link.



- The end host sends a packet to FW1.
- FW1 examines the contents of the packet to match it to an existing session. If there is no session match, FW1 determines that it has received the first packet for a new session and therefore becomes the session owner (assuming **Session Owner Selection** is set to **First Packet**).
- FW1 uses the configured session setup load-sharing option to identify the session setup firewall. In this example, FW2 is configured to perform session setup.
- FW1 uses the HA3 link to send the first packet to FW2.
- FW2 sets up the session and returns the packet to FW1 for Layer 7 processing, if any.
- FW1 then forwards the packet out the egress interface to the destination.

The following figure and text describe the path of a packet that matches an existing session:

© 2018 Palo Alto Networks, Inc.



- The end host sends a packet to FW1.
- FW1 examines the contents of the packet to match it to an existing session. If the session matches an existing session, FW1 processes the packet and sends the packet out the egress interface to the destination.

NAT in Active/Active HA Mode

In an active/active HA configuration:

- You must bind each Dynamic IP (DIP) NAT rule and Dynamic IP and Port (DIPP) NAT rule to either Device ID 0 or Device ID 1.
- You must bind each static NAT rule to either Device ID 0, Device ID 1, both Device IDs, or the firewall in active-primary state.

Thus, when one of the firewalls creates a new session, the Device ID **0** or Device ID **1** binding determines which NAT rules match the firewall. The device binding must include the session owner firewall to produce a match.

The session setup firewall performs the NAT policy match, but the NAT rules are evaluated based on the session owner. That is, the session is translated according to NAT rules that are bound to the session owner firewall. While performing NAT policy matching, a firewall skips all NAT rules that are not bound to the session owner firewall.

For example, suppose the firewall with Device ID 1 is the session owner and session setup firewall. When the firewall with Device ID 1 tries to match a session to a NAT rule, it skips all rules bound to Device ID 0. The firewall performs the NAT translation only if the session owner and the Device ID in the NAT rule match.

You will typically create device-specific NAT rules when the peer firewalls use different IP addresses for translation.

If one of the peer firewalls fails, the active firewall continues to process traffic for synchronized sessions from the failed firewall, including NAT traffic. In a source NAT configuration, when one firewall fails:

- The floating IP address that is used as the Translated IP address of the NAT rule transfers to the surviving firewall. Hence, the existing sessions that fail over will still use this IP address.
- All new sessions will use the device-specific NAT rules that the surviving firewall naturally owns. That is, the surviving firewall translates new sessions using only the NAT rules that match its Device ID; it ignores any NAT rules bound to the failed Device ID.

For examples of active/active HA with NAT, see:

© 2018 Palo Alto Networks, Inc.

- Use Case: Configure Active/Active HA with Source DIPP NAT Using Floating IP Addresses
- Use Case: Configure Separate Source NAT IP Address Pools for Active/Active HA Firewalls
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3

ECMP in Active/Active HA Mode

When an active/active HA peer fails, its sessions transfer to the new active-primary firewall, which tries to use the same egress interface that the failed firewall was using. If the firewall finds that interface among the ECMP paths, the transferred sessions will take the same egress interface and path. This behavior occurs regardless of the ECMP algorithm in use; using the same interface is desirable.

Only if no ECMP path matches the original egress interface will the active-primary firewall select a new ECMP path.

If you did not configure the same interfaces on the active/active peers, upon failover the active-primary firewall selects the next best path from the FIB table. Consequently, the existing sessions might not be distributed according to the ECMP algorithm.

© 2018 Palo Alto Networks, Inc.

Set Up Active/Passive HA

- Prerequisites for Active/Passive HA on page 267
- Configuration Guidelines for Active/Passive HA on page 267
- Configure Active/Passive HA on page 269
- Define HA Failover Conditions on page 274
- Verify Failover on page 275

Prerequisites for Active/Passive HA

To set up high availability on your Palo Alto Networks firewalls, you need a pair of firewalls that meet the following requirements:

- **The same model**—Both the firewalls in the pair must be of the same hardware model or virtual machine model.
- **The same PAN-OS version**—Both the firewalls should be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases.
- The same multi virtual system capability—Both firewalls must have **Multi Virtual System Capability** either enabled or not enabled. When enabled, each firewall requires its own multiple virtual systems licenses.
- **The same type of interfaces**—Dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
  - Determine the IP address for the HA1 (control) connection between the HA peers. The HA1 IP address for both peers must be on the same subnet if they are directly connected or are connected to the same switch.

    For firewalls without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both firewalls. However, because the management ports will not be directly cabled between the peers, make sure that you have a route that connects these two interfaces across your network.
  - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use Layer 3 only if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
- **The same set of licenses**—Licenses are unique to each firewall and cannot be shared between the firewalls. Therefore, you must license both firewalls identically. If both firewalls do not have an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.



As a best practice, if you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration Reset the Firewall to Factory Default Settings *on the new firewall. This ensures that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary firewall to the newly introduced firewall with the clean configuration.*

Configuration Guidelines for Active/Passive HA

To set up an active (PeerA) passive (PeerB) pair in HA, you must configure some options identically on both firewalls and some independently (non-matching) on each firewall. These HA settings are not synchronized between the firewalls. For details on what is/is not synchronized, see Reference: HA Synchronization on page 297.

© 2018 Palo Alto Networks, Inc.

The following checklist details the settings that you must configure identically on both firewalls:

- ❑ You must enable HA on both firewalls.
- ❑ You must configure the same Group ID value on both firewalls. The firewall uses the Group ID value to create a virtual MAC address for all the configured interfaces. See Floating IP Address and Virtual MAC Address for information about virtual MAC addresses. When a new active firewall takes over, it sends Gratuitous ARP messages from each of its connected interfaces to inform the connected Layer 2 switches of the virtual MAC address' new location.
- ❑ If you are using in-band ports as HA links, you must set the interfaces for the HA1 and HA2 links to type HA.
- ❑ Set the HA Mode to Active Passive on both firewalls.
- ❑ If required, enable preemption on both firewalls. The device priority value, however, must not be identical.
- ❑ If required, configure encryption on the HA1 link (for communication between the HA peers) on both firewalls.
- ❑ Based on the combination of HA1 and HA1 Backup ports you are using, use the following recommendations to decide whether you should enable heartbeat backup:

  HA functionality (HA1 and HA1 backup) is not supported on the management interface if it's configured for DHCP addressing (IP Type set to DHCP Client), except for AWS.

  - HA1: Dedicated HA1 port

    HA1 Backup: In-band port

    Recommendation: Enable Heartbeat Backup
  - HA1: Dedicated HA1 port

    HA1 Backup: Management port

    Recommendation: Do not enable Heartbeat Backup
  - HA1: In-band port

    HA1 Backup: In-band port

    Recommendation: Enable Heartbeat Backup
  - HA1: Management port

    HA1 Backup: In-band port

    Recommendation: Do not enable Heartbeat Backup

The following table lists the HA settings that you must configure independently on each firewall. See Reference: HA Synchronization on page 297 for more information about other configuration settings are not automatically synchronized between peers.

Independent Configuration Settings	PeerA	PeerB
Control Link	IP address of the HA1 link configured on this firewall (PeerA).	IP address of the HA1 link configured on this firewall (PeerB).
	For firewalls without dedicated HA ports, use the management port IP address for the control link.	
Data Link	By default, the HA2 link uses Ethernet/ Layer 2.	By default, the HA2 link uses Ethernet/Layer 2.

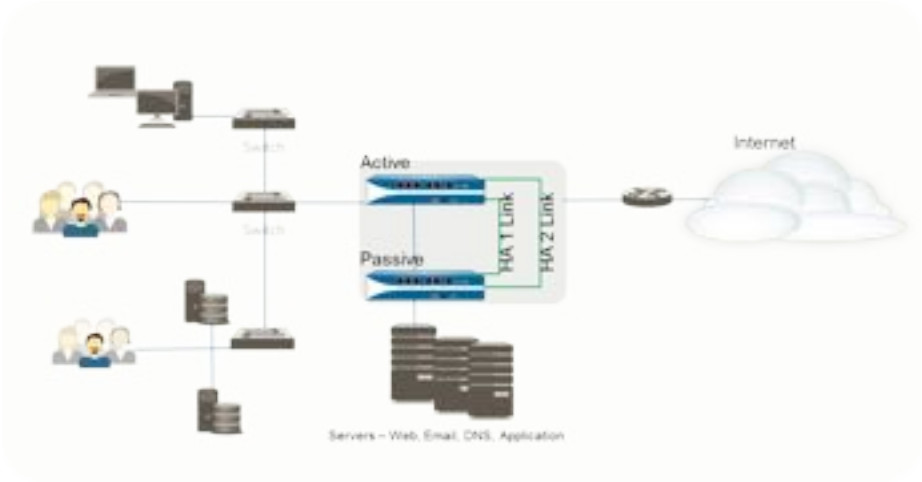
© 2018 Palo Alto Networks, Inc.

Independent Configuration Settings	PeerA	PeerB
The data link information is synchronized between the firewalls after HA is enabled and the control link is established between the firewalls.	If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerA).	If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerB).
Device Priority (required, if preemption is enabled)	The firewall you plan to make active must have a lower numerical value than its peer. So, if Peer A is to function as the active firewall, keep the default value of 100 and increment the value on PeerB. If the firewalls have the same device priority value, they use the MAC address of their HA1 as the tie-breaker.	If PeerB is passive, set the device priority value to a number larger than the setting on PeerA. For example, set the value to 110.
Link Monitoring —Monitor one or more physical interfaces that handle vital traffic on this firewall and define the failure condition.	Select the physical interfaces on the firewall that you would like to monitor and define the failure condition (all or any) to trigger a failover.	Pick a similar set of physical interfaces that you would like to monitor on this firewall and define the failure condition (all or any) to trigger a failover.
Path Monitoring —Monitor one or more destination IP addresses that the firewall can use ICMP pings to ascertain responsiveness.	Define the failure condition (all or any), ping interval and the ping count. This is particularly useful for monitoring the availability of other interconnected networking devices. For example, monitor the availability of a router that connects to a server, connectivity to the server itself, or some other vital device that is in the flow of traffic. Make sure that the node/device that you are monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a a path monitoring failure and trigger a failover.	Pick a similar set of devices or destination IP addresses that can be monitored for determining the failover trigger for PeerB. Define the failure condition (all or any), ping interval and the ping count.

Configure Active/Passive HA

The following procedure shows how to configure a pair of firewalls in an active/passive deployment as depicted in the following example topology.

© 2018 Palo Alto Networks, Inc.



To configure an active/passive HA pair, first complete the following workflow on the first firewall and then repeat the steps on the second firewall.

STEP 1 | Connect the HA ports to set up a physical connection between the firewalls.

- For firewalls with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on peers. Use a crossover cable if the peers are directly connected to each other.
- For firewalls without dedicated HA ports, select two data interfaces for the HA2 link and the backup HA1 link. Then, use an Ethernet cable to connect these in-band HA interfaces across both firewalls.

Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network.

STEP 2 | Enable ping on the management port.

Enabling ping allows the management port to exchange heartbeat backup information.

1. Select **Device** > **Setup** > **Management** and edit the Management Interface Settings.
2. Select **Ping** as a service that is permitted on the interface.

STEP 3 | If the firewall does not have dedicated HA ports, set up the data ports to function as HA ports.

For firewalls with dedicated HA ports continue to the next step.

1. Select **Network** > **Interfaces**.
2. Confirm that the link is up on the ports that you want to use.
3. Select the interface and set **Interface Type** to **HA**.
4. Set the **Link Speed** and **Link Duplex** settings, as appropriate.

STEP 4 | Set the HA mode and group ID.

1. Select **Device** > **High Availability** > **General** and edit the Setup section.
2. Set a **Group ID** and optionally a **Description** for the pair. The Group ID uniquely identifies each HA pair on your network. If you have multiple HA pairs that share the same broadcast domain you must set a unique Group ID for each pair.
3. Set the mode to **Active Passive**.

STEP 5 | Set up the control link connection.

This example shows an in-band port that is set to interface type HA.

For firewalls that use the management port as the control link, the IP address information is automatically pre-populated.

© 2018 Palo Alto Networks, Inc.

1. In **Device** > **High Availability** > **General**, edit the Control Link (HA1) section.
2. Select the **Port** that you have cabled for use as the HA1 link.
3. Set the **IPv4/IPv6 Address** and **Netmask**.

   If the HA1 interfaces are on separate subnets, enter the IP address of the **Gateway**. Do not add a gateway address if the firewalls are directly connected

STEP 6 | (Optional) Enable encryption for the control link connection.

This is typically used to secure the link if the two firewalls are not directly connected, that is if the ports are connected to a switch or a router.

1. Export the HA key from one firewall and import it into the peer firewall.
   1. Select **Device** > **Certificate Management** > **Certificates**.
   2. Select **Export HA key**. Save the HA key to a network location that the peer can access.
   3. On the peer firewall, select **Device** > **Certificate Management** > **Certificates**, and select **Import HA key** to browse to the location that you saved the key and import it in to the peer.
2. Select **Device** > **High Availability** > **General**, edit the Control Link (HA1) section.
3. Select **Encryption Enabled**.

STEP 7 | Set up the backup control link connection.

1. In **Device** > **High Availability** > **General**, edit the Control Link (HA1 Backup) section.
2. Select the HA1 backup interface and set the **IPv4/IPv6 Address** and **Netmask**.

STEP 8 | Set up the data link connection (HA2) and the backup HA2 connection between the firewalls.

1. In **Device** > **High Availability** > **General**, edit the Data Link (HA2) section.
2. Select the **Port** to use for the data link connection.
3. Select the **Transport** method. The default is **ethernet**, and will work when the HA pair is connected directly or through a switch. If you need to route the data link traffic through the network, select **IP** or **UDP** as the transport mode.
4. If you use IP or UDP as the transport method, enter the **IPv4/IPv6 Address** and **Netmask**.
5. Verify that **Enable Session Synchronization** is selected.
6. Select **HA2 Keep-alive** to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set (default is 10000 ms), the defined action will occur. For active/passive configuration, a critical system log message is generated when an HA2 keep-alive failure occurs.

   You can configure the HA2 keep-alive option on both firewalls, or just one firewall in the HA pair. If the option is only enabled on one firewall, only that firewall will send the keep-alive messages. The other firewall will be notified if a failure occurs.
7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**.

STEP 9 | Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port.

You do not need to enable heartbeat backup if you are using the management port for the control link.

1. In **Device** > **High Availability** > **General**, edit the Election Settings.
2. Select **Heartbeat Backup**.

   To allow the heartbeats to be transmitted between the firewalls, you must verify that the management port across both peers can route to each other.

© 2018 Palo Alto Networks, Inc.

Enabling heartbeat backup also allows you to prevent a split-brain situation. Split brain occurs when the HA1 link goes down causing the firewall to miss heartbeats, although the firewall is still functioning. In such a situation, each peer believes that the other is down and attempts to start services that are running, thereby causing a split brain. When the heartbeat backup link is enabled, split brain is prevented because redundant heartbeats and hello messages are transmitted over the management port.

STEP 10 | Set the device priority and enable preemption.

This setting is only required if you wish to make sure that a specific firewall is the preferred active firewall. For information, see Device Priority and Preemption.

1. In **Device** > **High Availability** > **General**, edit the Election Settings.
2. Set the numerical value in **Device Priority**. Make sure to set a lower numerical value on the firewall that you want to assign a higher priority to.

   If both firewalls have the same device priority value, the firewall with the lowest MAC address on the HA1 control link will become the active firewall.
3. Select **Preemptive**.

   You must enable preemptive on both the active firewall and the passive firewall.

STEP 11 | (Optional) Modify the HA Timers.

By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments.

1. In **Device** > **High Availability** > **General**, edit the Election Settings.
2. Select the **Aggressive** profile for triggering failover faster; select **Advanced** to define custom values for triggering failover in your set up.

   To view the preset value for an individual timer included in a profile, select Advanced and click Load Recommended or Load Aggressive. The preset values for your hardware model will be displayed on screen.

STEP 12 | (Optional), only configured on the passive firewall) Modify the link status of the HA ports on the passive firewall.

The passive link state is shutdown, by default. After you enable HA, the link state for the HA ports on the active firewall will be green and those on the passive firewall will be down and display as red.

Setting the link state to **Auto** allows for reducing the amount of time it takes for the passive firewall to take over when a failover occurs and it allows you to monitor the link state.

To enable the link status on the passive firewall to stay up and reflect the cabling status on the physical interface:

1. In **Device** > **High Availability** > **General**, edit the Active Passive Settings.
2. Set the **Passive Link State** to **Auto**.

   The auto option decreases the amount of time it takes for the passive firewall to take over when a failover occurs.

   Although the interface displays green (as cabled and up) it continues to discard all traffic until a failover is triggered.

© 2018 Palo Alto Networks, Inc.

When you modify the passive link state, make sure that the adjacent devices do not forward traffic to the passive firewall based only on the link status of the firewall.

STEP 13 | Enable HA.

1. Select **Device** > **High Availability** > **General** and edit the Setup section.
2. Select **Enable HA**.
3. Select **Enable Config Sync**. This setting enables the synchronization of the configuration settings between the active and the passive firewall.
4. Enter the IP address assigned to the control link of the peer in **Peer HA1 IP Address**.

   For firewalls without dedicated HA ports, if the peer uses the management port for the HA1 link, enter the management port IP address of the peer.
5. Enter the **Backup HA1 IP Address**.

STEP 14 | (Optional) Enable LACP and LLDP Pre-Negotiation for Active/Passive HA for faster failover if your network uses LACP or LLDP.

Enable LACP and LLDP before configuring HA pre-negotiation for the protocol if you want pre-negotiation to function in active mode.

1. Ensure that in Step 12 you set the link state to **Auto**.
2. Select **Network** > **Interfaces** > **Ethernet**.
3. To enable LACP active pre-negotiation:
   1. Select an AE interface in a Layer 2 or Layer 3 deployment.
   2. Select the **LACP** tab.
   3. Select **Enable in HA Passive State**.
   4. Click **OK**.

   You cannot also select Same System MAC Address for Active-Passive HA because pre-negotiation requires unique interface MAC addresses on the active and passive firewalls.
4. To enable LACP passive pre-negotiation:
   1. Select an Ethernet interface in a virtual wire deployment.
   2. Select the **Advanced** tab.
   3. Select the **LACP** tab.
   4. Select **Enable in HA Passive State**.
   5. Click **OK**.
5. To enable LLDP active pre-negotiation:
   1. Select an Ethernet interface in a Layer 2, Layer 3, or virtual wire deployment.
   2. Select the **Advanced** tab.
   3. Select the **LLDP** tab.
   4. Select **Enable in HA Passive State**.
   5. Click **OK**.

   If you want to allow LLDP passive pre-negotiation for a virtual wire deployment, perform Step e but do not enable LLDP itself.

STEP 15 | Save your configuration changes.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

STEP 16 | After you finish configuring both firewalls, verify that the firewalls are paired in active/passive HA.

1. Access the **Dashboard** on both firewalls, and view the High Availability widget.
2. On the active firewall, click the **Sync to peer** link.
3. Confirm that the firewalls are paired and synced, as shown as follows:
   - On the passive firewall: the state of the local firewall should display **passive** and the Running Config should show as **synchronized**.
   - On the active firewall: The state of the local firewall should display **active** and the Running Config should show as **synchronized**.

Define HA Failover Conditions

Perform the following task to define failover conditions and thus establish what will cause a firewall in an HA pair to fail over, an event where the task of securing traffic passes from the previously active firewall to its HA peer. The HA Overview describes conditions that cause a failover.



If you are using SNMPv3 to monitor the firewalls, note that the SNMPv3 Engine ID is unique to each firewall; the EngineID is not synchronized between the HA pair and, therefore, allows you to independently monitor each firewall in the HA pair. For information on setting up SNMP, see Forward Traps to an SNMP Manager*. Because the EngineID is generated using the firewall serial number, on the VM-Series firewall you must apply a valid license in order to obtain a unique EngineID for each firewall.*

STEP 1 | To configure link monitoring, define the interfaces you want to monitor. A change in the link state of these interfaces will trigger a failover.

1. Select **Device** > **High Availability** > **Link and Path Monitoring** and **Add** a Link Group.
2. Name the **Link Group**, **Add** the interfaces to monitor, and select the **Failure Condition** for the group. The Link group you define is added to the **Link Group** section.

STEP 2 | (Optional) Modify the failure condition for the Link Groups that you configured (in the preceding step) on the firewall.

By default, the firewall will trigger a failover when any monitored link fails.

1. Select the **Link Monitoring** section.
2. Set the **Failure Condition** to **All**.

   The default setting is **Any**.

STEP 3 | To configure path monitoring, define the destination IP addresses that the firewall should ping to verify network connectivity.

1. In the **Path Group** section of the **Device** > **High Availability** > **Link and Path Monitoring** tab, pick the **Add option for your set up:** Virtual Wire, VLAN, or Virtual Router.
2. Select the appropriate item from the drop-down for the **Name** and **Add** the IP addresses (source and/or destination, as prompted) that you wish to monitor. Then select the **Failure Condition** for the group. The path group you define is added to the **Path Group** section.

STEP 4 | (Optional) Modify the failure condition for all Path Groups configured on the firewall.

By default, the firewall will trigger a failover when any monitored path fails.

Set the **Failure Condition** to **All**.

The default setting is **Any**.

© 2018 Palo Alto Networks, Inc.

STEP 5 | **Commit** the configuration.

Verify Failover

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

STEP 1 | Suspend the active firewall.

Select **Device** > **High Availability** > **Operational Commands** and click the **Suspend local device** link.

STEP 2 | Verify that the passive firewall has taken over as active.

On the **Dashboard**, verify that the state of the passive firewall changes to **active** in the High Availability widget.

STEP 3 | Restore the suspended firewall to a functional state. Wait for a couple of minutes, and then verify that preemption has occurred, if **Preemptive** is enabled.

1. On the firewall you previously suspended, select **Device** > **High Availability** > **Operational Commands** and click the **Make local device functional** link.
2. In the High Availability widget on the **Dashboard**, confirm that the firewall has taken over as the active firewall and that the peer is now in a passive state.

© 2018 Palo Alto Networks, Inc.

Set Up Active/Active HA

- Prerequisites for Active/Active HA on page 276
- Configure Active/Active HA on page 277
- Determine Your Active/Active Use Case on page 281

Prerequisites for Active/Active HA

To set up active/active HA on your firewalls, you need a pair of firewalls that meet the following requirements:

- **The same model**—The firewalls in the pair must be of the same hardware model.
- **The same PAN-OS version**—The firewalls must be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases.
- **The same multi virtual system capability**—Both firewalls must have **Multi Virtual System Capability** either enabled or not enabled. When enabled, each firewall requires its own multiple virtual systems licenses.
- **The same type of interfaces**—Dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
  - The HA interfaces must be configured with static IP addresses only, not IP addresses obtained from DHCP (except AWS can use DHCP addresses). Determine the IP address for the HA1 (control) connection between the HA peers. The HA1 IP address for the peers must be on the same subnet if they are directly connected or are connected to the same switch.

    For firewalls without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both firewalls. However, because the management ports will not be directly cabled between the peers, make sure that you have a route that connects these two interfaces across your network.
  - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use Layer 3 only if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
  - Each firewall needs a dedicated interface for the HA3 link. The PA-7000 Series firewalls use the HSCI port for HA3. The PA-5200 Series firewalls can use the HSCI port for HA3 or you can configure aggregate interfaces on the dataplane ports for HA3 for redundancy. On the remaining platforms, you can configure aggregate interfaces on dataplane ports as the HA3 link for redundancy.
- **The same set of licenses**—Licenses are unique to each firewall and cannot be shared between the firewalls. Therefore, you must license both firewalls identically. If both firewalls do not have an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.



If you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration, it is recommended that you Reset the Firewall to Factory Default Settings *on the new firewall. This will ensure that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary firewall to the newly introduced firewall with the clean config. You will also have to configure local IP addresses.*

© 2018 Palo Alto Networks, Inc.

Configure Active/Active HA

The following procedure describes the basic workflow for configuring your firewalls in an active/active configuration. However, before you begin, Determine Your Active/Active Use Case for configuration examples more tailored to your specific network environment.

If you have a switch located between your HA firewalls, the switch ports that connect the HA3 link must support jumbo frames to handle the overhead associated with the MAC-in-MAC encapsulation on the HA3 link.

To configure active/active, first complete the following steps on one peer and then complete them on the second peer, ensuring that you set the Device ID to different values (0 or 1) on each peer.

STEP 1 | Connect the HA ports to set up a physical connection between the firewalls.

For each use case, the firewalls could be any hardware model; choose the HA3 step that corresponds with your model.

- For firewalls with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on peers. Use a crossover cable if the peers are directly connected to each other.
- For firewalls without dedicated HA ports, select two data interfaces for the HA2 link and the backup HA1 link. Then, use an Ethernet cable to connect these in-band HA interfaces across both firewalls. Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network.
- For HA3:
  - On PA-7000 Series firewalls, connect the High Speed Chassis Interconnect (HSCI-A) on the first chassis to the HSCI-A on the second chassis, and the HSCI-B on the first chassis to the HSCI-B on the second chassis. On a PA-5200 Series firewall (which has one HSCI port), connect the HSCI port on the first chassis to the HSCI port on the second chassis. You can also use data ports for HA3 on PA-5200 Series firewalls.
  - On any other hardware model, use dataplane interfaces for HA3.

STEP 2 | Enable ping on the management port.

Enabling ping allows the management port to exchange heartbeat backup information.

1. In **Device** > **Setup** > **Management**, edit Management Interface Settings.
2. Select **Ping** as a service that is permitted on the interface.

STEP 3 | If the firewall does not have dedicated HA ports, set up the data ports to function as HA ports.

For firewalls with dedicated HA ports continue to the next step.

1. Select **Network** > **Interfaces**.
2. Confirm that the link is up on the ports that you want to use.
3. Select the interface and set **Interface Type** to **HA**.
4. Set the **Link Speed** and **Link Duplex** settings, as appropriate.

STEP 4 | Enable active/active HA and set the group ID.

1. In **Device** > **High Availability** > **General**, edit Setup.
2. Select **Enable HA**.
3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63).
4. (Optional) Enter a **Description**.

© 2018 Palo Alto Networks, Inc.

5. For **Mode**, select **Active Active**.

STEP 5 | Set the Device ID, enable synchronization, and identify the control link on the peer firewall

1. In **Device** > **High Availability** > **General**, edit Setup.
2. Select **Device ID** as follows:
   - When configuring the first peer, set the **Device ID** to **0**.
   - When configuring the second peer, set the **Device ID** to **1**.
3. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default).
4. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall.
5. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall.
6. Click **OK**.

STEP 6 | Determine whether or not the firewall with the lower Device ID preempts the active-primary firewall upon recovery from a failure.

1. In **Device** > **High Availability** > **General**, edit Election Settings.
2. Select **Preemptive** to cause the firewall with the lower Device ID to automatically resume active-primary operation after either firewall recovers from a failure. Both firewalls must have **Preemptive** selected for preemption to occur.

   Leave **Preemptive** unselected if you want the active-primary role to remain with the current firewall until you manually make the recovered firewall the active-primary firewall.

STEP 7 | Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port.

You need not enable heartbeat backup if you are using the management port for the control link.

1. In **Device** > **High Availability** > **General**, edit Election Settings.
2. Select **Heartbeat Backup**.

   To allow the heartbeats to be transmitted between the firewalls, you must verify that the management port across both peers can route to each other.



Enabling heartbeat backup allows you to prevent a split-brain situation. Split brain occurs when the HA1 link goes down, causing the firewall to miss heartbeats, although the firewall is still functioning. In such a situation, each peer believes the other is down and attempts to start services that are running, thereby causing a split brain. Enabling heartbeat backup prevents split brain because redundant heartbeats and hello messages are transmitted over the management port.

STEP 8 | (Optional) Modify the HA Timers.

By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments.

1. In **Device** > **High Availability** > **General**, edit Election Settings.
2. Select **Aggressive** to trigger faster failover. Select **Advanced** to define custom values for triggering failover in your setup.



To view the preset value for an individual timer included in a profile, select Advanced and click Load Recommended or Load Aggressive. The preset values for your hardware model will be displayed on screen.

STEP 9 | Set up the control link connection.

© 2018 Palo Alto Networks, Inc.

This example uses an in-band port that is set to interface type HA.

For firewalls that use the management port as the control link, the IP address information is automatically pre-populated.

1. In **Device** > **High Availability** > **General**, edit Control Link (HA1).
2. Select the **Port** that you have cabled for use as the HA1 link.
3. Set the **IPv4/IPv6 Address** and **Netmask**.

   If the HA1 interfaces are on separate subnets, enter the IP address of the **Gateway**. Do not add a gateway address if the firewalls are directly connected.

STEP 10 | (Optional) Enable encryption for the control link connection.

This is typically used to secure the link if the two firewalls are not directly connected, that is if the ports are connected to a switch or a router.

1. Export the HA key from one firewall and import it into the peer firewall.
   1. Select **Device** > **Certificate Management** > **Certificates**.
   2. Select **Export HA key**. Save the HA key to a network location that the peer can access.
   3. On the peer firewall, select **Device** > **Certificate Management** > **Certificates**, and select **Import HA key** to browse to the location that you saved the key and import it in to the peer.
2. In **Device** > **High Availability** > **General**, edit the Control Link (HA1).
3. Select **Encryption Enabled**.

STEP 11 | Set up the backup control link connection.

1. In **Device** > **High Availability** > **General**, edit Control Link (HA1 Backup).
2. Select the HA1 backup interface and set the **IPv4/IPv6 Address** and **Netmask**.

STEP 12 | Set up the data link connection (HA2) and the backup HA2 connection between the firewalls.

1. In **Device** > **High Availability** > **General**, edit Data Link (HA2).
2. Select the **Port** to use for the data link connection.
3. Select the **Transport** method. The default is **ethernet**, and will work when the HA pair is connected directly or through a switch. If you need to route the data link traffic through the network, select **IP** or **UDP** as the transport mode.
4. If you use IP or UDP as the transport method, enter the **IPv4/IPv6 Address** and **Netmask**.
5. Verify that **Enable Session Synchronization** is selected.
6. Select **HA2 Keep-alive** to enable monitoring on the HA2 data link between the HA peers. If a failure occurs based on the threshold that is set (default is 10000 ms), the defined action will occur. For active/passive configuration, a critical system log message is generated when an HA2 keep-alive failure occurs.

   You can configure the HA2 keep-alive option on both firewalls, or just one firewall in the HA pair. If the option is only enabled on one firewall, only that firewall will send the keep-alive messages. The other firewall will be notified if a failure occurs.
7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**.
8. Click **OK**.

STEP 13 | Configure the HA3 link for packet forwarding.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **HA3 Interface**, select the interface you want to use to forward packets between active/active HA peers. It must be a dedicated interface capable of Layer 2 transport and set to **Interface Type HA**.

© 2018 Palo Alto Networks, Inc.

3. Select **VR Sync** to force synchronization of all virtual routers configured on the HA peers. Select when the virtual router is not configured for dynamic routing protocols. Both peers must be connected to the same next-hop router through a switched network and must use static routing only.
4. Select **QoS Sync** to synchronize the QoS profile selection on all physical interfaces. Select when both peers have similar link speeds and require the same QoS profiles on all physical interfaces. This setting affects the synchronization of QoS settings on the **Network** tab. QoS policy is synchronized regardless of this setting.

STEP 14 | (Optional) Modify the Tentative Hold time.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **Tentative Hold Time (sec)**, enter the number of seconds that a firewall stays in Tentative state after it fails (range is 10-600, default is 60).

STEP 15 | Configure Session Owner and Session Setup.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **Session Owner Selection**, select one of the following:
   - **First Packet**—The firewall that receives the first packet of a new session is the session owner (recommended setting). This setting minimizes traffic across HA3 and load shares traffic across peers.
   - **Primary Device**—The firewall that is in active-primary state is the session owner.
3. For **Session Setup**, select one of the following:
   - **IP Modulo**— Distributes session setup load based on parity of the source IP address.
   - **Primary Device**—The active-primary firewall sets up all sessions.
   - **First Packet**—The firewall that receives the first packet of a new session performs session setup (recommended setting).

     Start with First Packet for Session Owner and Session Setup, and then based on load distribution, you can change to one of the other options.
   - **IP Hash**—The firewall uses a hash of either the source IP address or a combination of the source and destination IP addresses to distribute session setup responsibilities.
4. Click **OK**.

STEP 16 | Configure an HA virtual address.

You need a virtual address to use a Floating IP Address and Virtual MAC Address or ARP Load-Sharing.

1. In **Device** > **High Availability** > **Active/Active Config**, **Add** a Virtual Address.
2. Enter or select an **Interface**.
3. Select the **IPv4** or **IPv6** tab and click **Add**.
4. Enter an **IPv4 Address** or **IPv6 Address**.
5. For **Type**:
   - Select **Floating** to configure the virtual IP address to be a floating IP address.
   - Select **ARP Load Sharing** to configure the virtual IP address to be a shared IP address and skip to Configure ARP Load-Sharing.

STEP 17 | Configure the floating IP address.

1. Do not select **Floating IP bound to the Active-Primary device** unless you want the active/active HA pair to behave like an active/passive HA pair.
2. For **Device 0 Priority** and **Device 1 Priority**, enter a priority for the firewall configured with Device ID 0 and Device ID 1, respectively. The relative priorities determine which peer owns the floating

© 2018 Palo Alto Networks, Inc.

IP address you just configured (range is 0-255). The firewall with the lowest priority value (highest priority) owns the floating IP address.

3. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down.
4. Click **OK**.

STEP 18 | Configure ARP Load-Sharing.

The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.

1. For **Device Selection Algorithm**, select one of the following:
   - **IP Modulo**—The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address.
   - **IP Hash**—The firewall that will respond to ARP requests is based on a hash of the ARP requester's IP address.
2. Click **OK**.

STEP 19 | Define HA Failover Conditions.

STEP 20 | **Commit** the configuration.

STEP 21 | Reboot the firewall after changing the jumbo frame configuration.

1. Select **Device** > **Setup** > **Operations**.
2. Click **Reboot Device**.

Determine Your Active/Active Use Case

Determine which type of use case you have and then select the corresponding procedure to configure active/active HA.

If you are using Route-Based Redundancy on page 260, Floating IP Address and Virtual MAC Address on page 257, or ARP Load-Sharing on page 259, select the corresponding procedure:

- Use Case: Configure Active/Active HA with Route-Based Redundancy on page 282
- Use Case: Configure Active/Active HA with Floating IP Addresses on page 282
- Use Case: Configure Active/Active HA with ARP Load-Sharing on page 284

If you want a Layer 3 active/active HA deployment that behaves like an active/passive deployment, select the following procedure:

- Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall on page 285

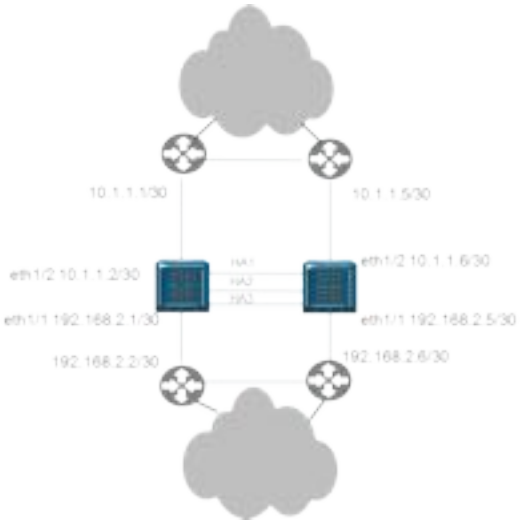
If you are configuring NAT in Active/Active HA Mode on page 265, see the following procedures:

- Use Case: Configure Active/Active HA with Source DIPP NAT Using Floating IP Addresses on page 287
- Use Case: Configure Separate Source NAT IP Address Pools for Active/Active HA Firewalls on page 290
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT on page 291
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3 on page 292

© 2018 Palo Alto Networks, Inc.

Use Case: Configure Active/Active HA with Route-Based Redundancy

The following Layer 3 topology illustrates two PA-7050 firewalls in an active/active HA environment that use Route-Based Redundancy. The firewalls belong to an OSPF area. When a link or firewall fails, OSPF handles the redundancy by redirecting traffic to the functioning firewall.



STEP 1 | Configure Active/Active HA.

Perform Step 1 through Step 15.

STEP 2 | Configure OSPF.

See OSPF.

STEP 3 | Define HA failover conditions.

Define HA Failover Conditions.

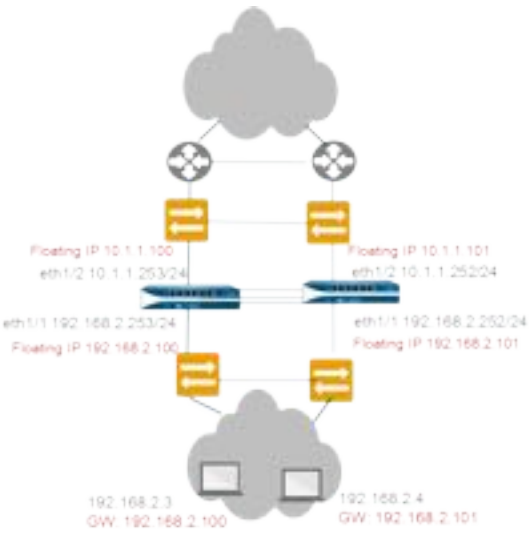
STEP 4 | **Commit** the configuration.

STEP 5 | Configure the peer firewall in the same way, except in Step 5, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.

Use Case: Configure Active/Active HA with Floating IP Addresses

In this Layer 3 interface example, the HA firewalls connect to switches and use floating IP addresses to handle link or firewall failures. The end hosts are each configured with a gateway, which is the floating IP address of one of the HA firewalls. See Floating IP Address and Virtual MAC Address.

© 2018 Palo Alto Networks, Inc.



STEP 1 | Configure Active/Active HA.

Perform Step 1 through Step 15.

STEP 2 | Configure an HA virtual address.

You need a virtual address to use a Floating IP Address and Virtual MAC Address.

1. In **Device** > **High Availability** > **Active/Active Config**, **Add** a Virtual Address.
2. Enter or select an **Interface**.
3. Select the **IPv4** or **IPv6** tab and click **Add**.
4. Enter an **IPv4 Address** or **IPv6 Address**.
5. For **Type**, select **Floating** to configure the virtual IP address to be a floating IP address.

STEP 3 | Configure the floating IP address.

1. Do not select **Floating IP bound to the Active-Primary device**.
2. For **Device 0 Priority** and **Device 1 Priority**, enter a priority for the firewall configured with Device ID 0 and Device ID 1, respectively. The relative priorities determine which peer owns the floating IP address you just configured (range is 0-255). The firewall with the lowest priority value (highest priority) owns the floating IP address.
3. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down.
4. Click **OK**.

STEP 4 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

Perform Step 19 of Configure Active/Active HA.

STEP 5 | Define HA Failover Conditions

STEP 6 | **Commit** the configuration.

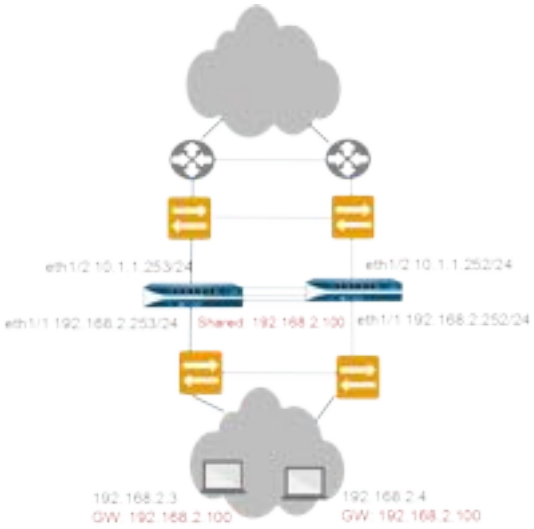
STEP 7 | Configure the peer firewall in the same way, except selecting a different Device ID.

For example, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.

© 2018 Palo Alto Networks, Inc.

Use Case: Configure Active/Active HA with ARP Load-Sharing

In this example, hosts in a Layer 3 deployment need gateway services from the HA firewalls. The firewalls are configured with a single shared IP address, which allows ARP Load-Sharing. The end hosts are configured with the same gateway, which is the shared IP address of the HA firewalls.



STEP 1 | Perform Step 1 through Step 15 of Configure Active/Active HA.

STEP 2 | Configure an HA virtual address.

The virtual address is the shared IP address that allows ARP Load-Sharing.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Enter or select an **Interface**.
3. Select the **IPv4** or **IPv6** tab and click **Add**.
4. Enter an **IPv4 Address** or **IPv6 Address**.
5. For **Type**, select **ARP Load Sharing**, which allows both peers to use the virtual IP address for ARP Load-Sharing.

STEP 3 | Configure ARP Load-Sharing.

The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.

1. For **Device Selection Algorithm**, select one of the following:
   - **IP Modulo**—The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address.
   - **IP Hash**—The firewall that will respond to ARP requests is based on a hash of the ARP requester's IP address.
2. Click **OK**.

STEP 4 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

STEP 5 | Define HA Failover Conditions

STEP 6 | **Commit** the configuration.

STEP 7 | Configure the peer firewall in the same way, except selecting a different Device ID.

© 2018 Palo Alto Networks, Inc.

For example, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.

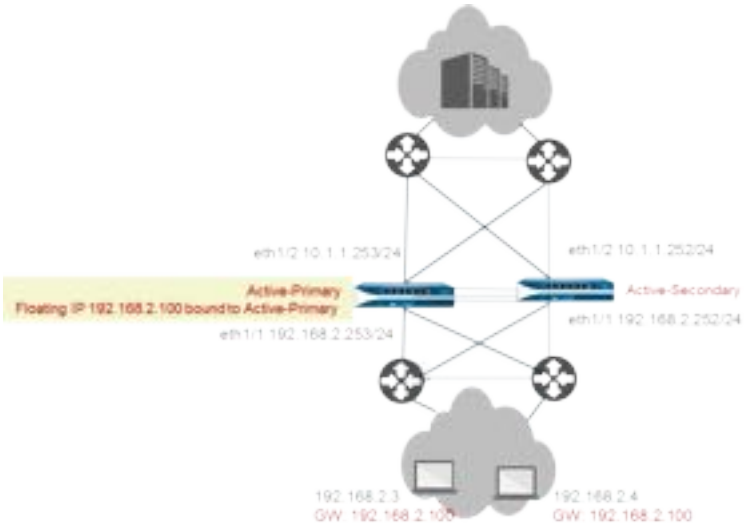
Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall

In mission-critical data centers, you may want both Layer 3 HA firewalls to participate in path monitoring so that they can detect path failures upstream from both firewalls. Additionally, you prefer to control if and when the floating IP address returns to the recovered firewall after it comes back up, rather than the floating IP address returning to the device ID to which it is bound. (That default behavior is described in Floating IP Address and Virtual MAC Address.)

In this use case, you control when the floating IP address and therefore the active-primary role move back to a recovered HA peer. The active/active HA firewalls share a single floating IP address that you bind to whichever firewall is in the active-primary state. With only one floating IP address, network traffic flows predominantly to a single firewall, so this active/active deployment functions like an active/passive deployment.

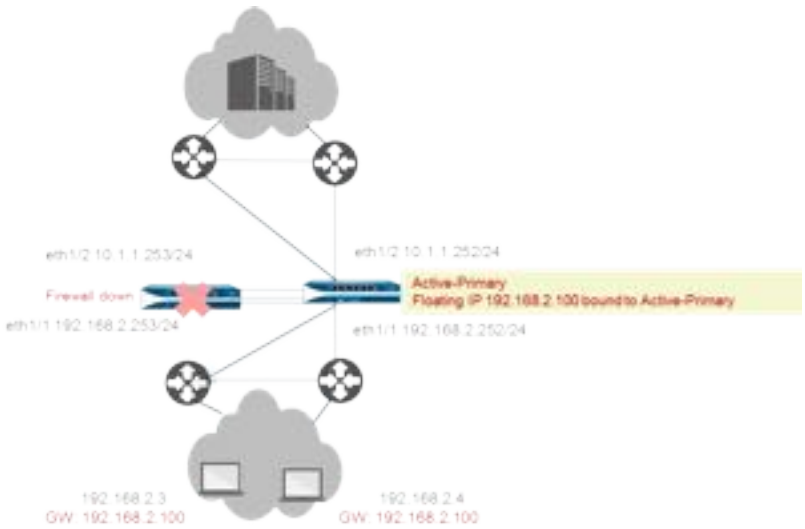
In this use case, Cisco Nexus 7010 switches with virtual PortChannels (vPCs) operating in Layer 3 connect to the firewalls. You must configure the Layer 3 switches (router peers) north and south of the firewalls with a route preference to the floating IP address. That is, you must design your network so the route tables of the router peers have the best path to the floating IP address. This example uses static routes with the proper metrics so that the route to the floating IP address uses a lower metric (the route to the floating IP address is preferred) and receives the traffic. An alternative to using static routes would be to design the network to redistribute the floating IP address into the OSPF routing protocol (if you are using OSPF).

The following topology illustrates the floating IP address bound to the active-primary firewall, which is initially Peer A, the firewall on the left.



Upon a failover, when the active-primary firewall (Peer A) goes down and the active-secondary firewall (Peer B) takes over as the active-primary peer, the floating IP address moves to Peer B (shown in the following figure). Peer B remains the active-primary firewall and traffic continues to go to Peer B, even when Peer A recovers and becomes the active-secondary firewall. You decide if and when to make Peer A the active-primary firewall again.

© 2018 Palo Alto Networks, Inc.



Binding the floating IP address to the active-primary firewall provides you with more control over how the firewalls determine floating IP address ownership as they move between various HA Firewall States. The following advantages result:

- You can have an active/active HA configuration for path monitoring out of both firewalls, but have the firewalls function like an active/passive HA configuration because traffic directed to the floating IP address always goes to the active-primary firewall.

When you disable preemption on both firewalls, you have the following additional benefits:

- The floating IP address does not move back and forth between HA firewalls if the active-secondary firewall flaps up and down.
- You can review the functionality of the recovered firewall and the adjacent components before manually directing traffic to it again, which you can do at a convenient down time.
- You have control over which firewall owns the floating IP address so that you keep all flows of new and existing sessions on the active-primary firewall, thereby minimizing traffic on the HA3 link.

- *We strongly recommended you configure HA link monitoring on the interface(s) that support the floating IP address(es) to allow each HA peer to quickly detect a link failure and fail over to its peer. Both HA peers must have link monitoring for it to function.*
- *We strongly recommend you configure HA path monitoring to notify each HA peer when a path has failed so a firewall can fail over to its peer. Because the floating IP address is always bound to the active-primary firewall, the firewall cannot automatically fail over to the peer when a path goes down and path monitoring is not enabled.*

You cannot configure NAT for a floating IP address that is bound to an active-primary firewall.

STEP 1 | Perform Step 1 through Step 5 of Configure Active/Active HA.

STEP 2 | (Optional) Disable preemption.

Disabling preemption allows you full control over when the recovered firewall becomes the active-primary firewall.

1. In **Device** > **High Availability** > **General**, edit the Election Settings.

© 2018 Palo Alto Networks, Inc.

2. Clear **Preemptive** if it is enabled.
3. Click **OK**.

STEP 3 | Perform Step 7 through Step 14 of Configure Active/Active HA.

STEP 4 | Configure Session Owner and Session Setup.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **Session Owner Selection**, we recommend you select **Primary Device**. The firewall that is in active-primary state is the session owner.

   Alternatively, for **Session Owner Selection** you can select **First Packet** and then for **Session Setup**, select **Primary Device** or **First Packet**.
3. For **Session Setup**, select **Primary Device**—The active-primary firewall sets up all sessions. This is the recommended setting if you want your active/active configuration to behave like an active/passive configuration because it keeps all activity on the active-primary firewall.

   You must also engineer your network to eliminate the possibility of asymmetric traffic going to the HA pair. If you don't do so and traffic goes to the active-secondary firewall, setting Session Owner Selection and Session Setup to Primary Device causes the traffic to traverse HA3 to get to the active-primary firewall for session ownership and session setup.
4. Click **OK**.

STEP 5 | Configure an HA virtual address.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Enter or select an **Interface**.
3. Select the **IPv4** or **IPv6** tab and **Add** an **IPv4 Address** or **IPv6 Address**.
4. For **Type**, select **Floating**, which configures the virtual IP address to be a floating IP address.
5. Click **OK**.

STEP 6 | Bind the floating IP address to the active-primary firewall.

1. Select **Floating IP bound to the Active-Primary device**.
2. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down.
3. Click **OK**.

STEP 7 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

STEP 8 | **Commit** the configuration.

STEP 9 | Configure the peer firewall in the same way, except selecting a different Device ID.

For example, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.

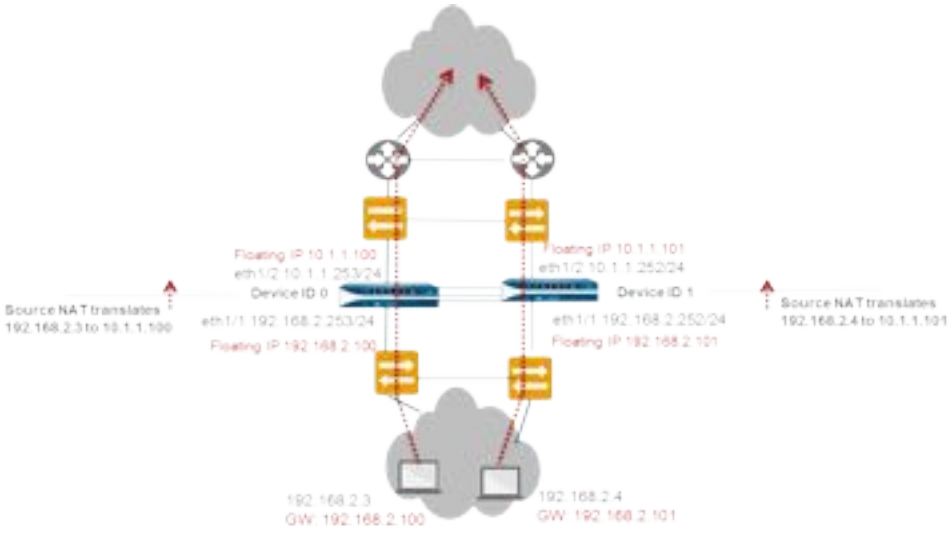
Use Case: Configure Active/Active HA with Source DIPP NAT Using Floating IP Addresses

This Layer 3 interface example uses source NAT in Active/Active HA Mode. The Layer 2 switches create broadcast domains to ensure users can reach everything north and south of the firewalls.

PA-3050-1 has Device ID 0 and its HA peer, PA-3050-2, has Device ID 1. In this use case, NAT translates the source IP address and port number to the floating IP address configured on the egress interface. Each host is configured with a default gateway address, which is the floating IP address on Ethernet1/1 of each

© 2018 Palo Alto Networks, Inc.

firewall. The configuration requires two source NAT rules, one bound to each Device ID, although you configure both NAT rules on a single firewall and they are synchronized to the peer firewall.



STEP 1 | On PA-3050-2 (Device ID 1), perform Step 1 through Step 3 of Configure Active/Active HA.

STEP 2 | Enable active/active HA.

1. In **Device** > **High Availability** > **General**, edit Setup.
2. Select **Enable HA**.
3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63).
4. For **Mode**, select **Active Active**.
5. Set the **Device ID** to **1**.
6. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default).
7. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall.
8. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall.
9. Click **OK**.

STEP 3 | Configure Active/Active HA.

Complete Step 6 through Step 14.

STEP 4 | Configure Session Owner and Session Setup.

1. In **Device** > **High Availability** > **Active/Active Config**, edit Packet Forwarding.
2. For **Session Owner Selection**, select **First Packet**—The firewall that receives the first packet of a new session is the session owner.
3. For **Session Setup**, select **IP Modulo**—Distributes session setup load based on parity of the source IP address.
4. Click **OK**.

STEP 5 | Configure an HA virtual address.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Select **Interface** eth1/1.
3. Select **IPv4** and **Add** an **IPv4 Address** of 10.1.1.101.

© 2018 Palo Alto Networks, Inc.

4. For **Type**, select **Floating**, which configures the virtual IP address to be a floating IP address.

STEP 6 | Configure the floating IP address.

1. Do not select **Floating IP bound to the Active-Primary device**.
2. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down.
3. Click **OK**.

STEP 7 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

STEP 8 | Define HA Failover Conditions.

STEP 9 | **Commit** the configuration.

STEP 10 | Configure the peer firewall, PA-3050-1 with the same settings, except for the following changes:

- Select **Device ID 0**.
- Configure an HA virtual address of 10.1.1.100.
- For **Device 1 Priority**, enter 255. For **Device 0 Priority**, enter 0.

In this example, Device ID 0 has a lower priority value so a higher priority; therefore, the firewall with Device ID 0 (PA-3050-1) owns the floating IP address 10.1.1.100.

STEP 11 | Still on PA-3050-1, create the source NAT rule for Device ID 0.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a **Name** for the rule that in this example identifies it as a source NAT rule for Device ID 0.
3. For **NAT Type**, select **ipv4** (default).
4. On the **Original Packet**, for **Source Zone**, select **Any**.
5. For **Destination Zone**, select the zone you created for the external network.
6. Allow **Destination Interface**, **Service**, **Source Address**, and **Destination Address** to remain set to **Any**.
7. For the **Translated Packet**, select **Dynamic IP And Port** for **Translation Type**.
8. For **Address Type**, select **Interface Address**, in which case the translated address will be the IP address of the interface. Select an **Interface** (eth1/1 in this example) and an **IP Address** of the floating IP address 10.1.1.100.
9. On the **Active/Active HA Binding** tab, for **Active/Active HA Binding**, select **0** to bind the NAT rule to Device ID 0.
10. Click **OK**.

STEP 12 | Create the source NAT rule for Device ID 1.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a **Name** for the policy rule that in this example helps identify it as a source NAT rule for Device ID 1.
3. For **NAT Type**, select **ipv4** (default).
4. On the **Original Packet**, for **Source Zone**, select **Any**. For **Destination Zone**, select the zone you created for the external network.
5. Allow **Destination Interface**, **Service**, **Source Address**, and **Destination Address** to remain set to **Any**.
6. For the **Translated Packet**, select **Dynamic IP And Port** for **Translation Type**.
7. For **Address Type**, select **Interface Address**, in which case the translated address will be the IP address of the interface. Select an **Interface** (eth1/1 in this example) and an **IP Address** of the floating IP address 10.1.1.101.

© 2018 Palo Alto Networks, Inc.

8. On the **Active/Active HA Binding** tab, for the **Active/Active HA Binding**, select **1** to bind the NAT rule to Device ID 1.
9. Click **OK**.

STEP 13 | **Commit** the configuration.

Use Case: Configure Separate Source NAT IP Address Pools for Active/Active HA Firewalls

If you want to use IP address pools for source NAT in Active/Active HA Mode on page 265, each firewall must have its own pool, which you then bind to a Device ID in a NAT rule.

Address objects and NAT rules are synchronized (in both active/passive and active/active mode), so they need to be configured on only one of the firewalls in the HA pair.

This example configures an address object named Dyn-IP-Pool-dev0 containing the IP address pool 10.1.1.140-10.1.1.150. It also configures an address object named Dyn-IP-Pool-dev1 containing the IP address pool 10.1.1.160-10.1.1.170. The first address object is bound to Device ID 0; the second address object is bound to Device ID 1.

STEP 1 | On one HA firewall, create address objects.

1. Select **Objects** > **Addresses** and **Add** an address object **Name**, in this example, Dyn-IP-Pool-dev0.
2. For **Type**, select **IP Range** and enter the range 10.1.1.140-10.1.1.150.
3. Click **OK**.
4. Repeat this step to configure another address object named Dyn-IP-Pool-dev1 with the **IP Range** of 10.1.1.160-10.1.1.170.

STEP 2 | Create the source NAT rule for Device ID 0.

1. Select **Policies** > **NAT** and **Add** a NAT policy rule with a **Name**, for example, Src-NAT-dev0.
2. For **Original Packet**, for **Source Zone**, select **Any**.
3. For **Destination Zone**, select the destination zone for which you want to translate the source address, such as Untrust.
4. For **Translated Packet**, for **Translation Type**, select **Dynamic IP and Port**.
5. For **Translated Address**, **Add** the address object you created for the pool of addresses belonging to Device ID 0: Dyn-IP-Pool-dev0.
6. For **Active/Active HA Binding**, select **0** to bind the NAT rule to Device ID 0.
7. Click **OK**.

STEP 3 | Create the source NAT rule for Device ID 1.

1. Select **Policies** > **NAT** and **Add** a NAT policy rule with a **Name**, for example, Src-NAT-dev1.
2. For **Original Packet**, for **Source Zone**, select **Any**.
3. For **Destination Zone**, select the destination zone for which you want to translate the source address, such as Untrust.
4. For **Translated Packet**, for **Translation Type**, select **Dynamic IP and Port**.
5. For **Translated Address**, **Add** the address object you created for the pool of addresses belonging to Device ID 1: Dyn-IP-Pool-dev1.
6. For **Active/Active HA Binding**, select **1** to bind the NAT rule to Device ID 1.
7. Click **OK**.

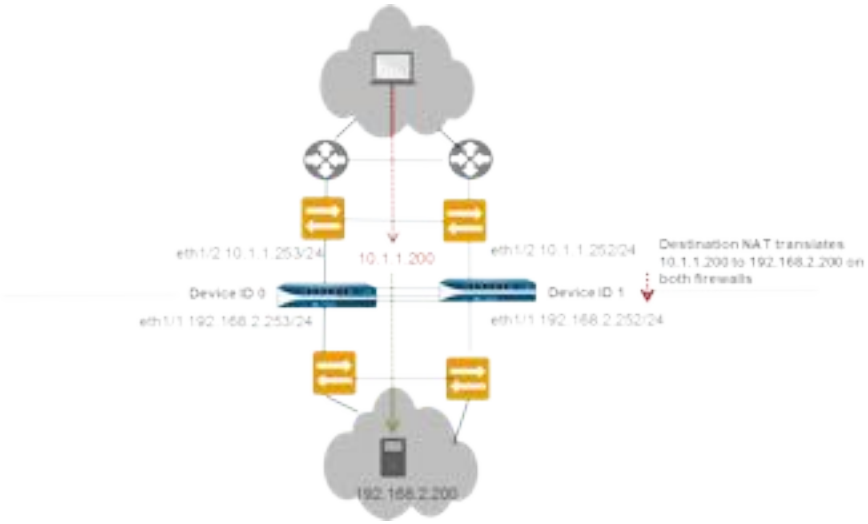
STEP 4 | **Commit** the configuration.

© 2018 Palo Alto Networks, Inc.

Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT

This Layer 3 interface example uses NAT in Active/Active HA Mode and ARP Load-Sharing with destination NAT. Both HA firewalls respond to an ARP request for the destination NAT address with the ingress interface MAC address. Destination NAT translates the public, shared IP address (in this example, 10.1.1.200) to the private IP address of the server (in this example, 192.168.2.200).

When the HA firewalls receive traffic for the destination 10.1.1.200, both firewalls could possibly respond to the ARP request, which could cause network instability. To avoid the potential issue, configure the firewall that is in active-primary state to respond to the ARP request by binding the destination NAT rule to the active-primary firewall.



STEP 1 | On PA-3050-2 (Device ID 1), perform Step 1 through Step 3 of Configure Active/Active HA.

STEP 2 | Enable active/active HA.

1. In **Device** > **High Availability** > **General**, edit Setup.
2. Select **Enable HA**.
3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63).
4. (Optional) Enter a **Description**.
5. For **Mode**, select **Active Active**.
6. Select **Device ID** to be **1**.
7. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default).
8. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall.
9. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall.
10. Click **OK**.

STEP 3 | Perform Step 6 through Step 15 in Configure Active/Active HA.

STEP 4 | Configure an HA virtual address.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Select **Interface** eth1/1.

© 2018 Palo Alto Networks, Inc.

3. Select **IPv4** and **Add** an **IPv4 Address** of 10.1.1.200.
4. For **Type**, select **ARP Load Sharing**, which configures the virtual IP address to be for both peers to use for ARP Load-Sharing.

STEP 5 | Configure ARP Load-Sharing.

The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.

1. For **Device Selection Algorithm**, select **IP Modulo**. The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address.
2. Click **OK**.

STEP 6 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

STEP 7 | Define HA Failover Conditions.

STEP 8 | **Commit** the configuration.

STEP 9 | Configure the peer firewall, PA-3050-1 (Device ID 0), with the same settings, except in Step 2 select **Device ID 0**.

STEP 10 | Still on PA-3050-1 (Device ID 0), create the destination NAT rule so that the active-primary firewall responds to ARP requests.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a **Name** for the rule that, in this example, identifies it as a destination NAT rule for Layer 2 ARP.
3. For **NAT Type**, select **ipv4** (default).
4. On the **Original Packet**, for **Source Zone**, select **Any**.
5. For **Destination Zone**, select the Untrust zone you created for the external network.
6. Allow **Destination Interface**, **Service**, and **Source Address** to remain set to **Any**.
7. For **Destination Address**, specify 10.1.1.200.
8. For the **Translated Packet**, Source Address Translation remains **None**.
9. For **Destination Address Translation**, enter the private IP address of the destination server, in this example, 192.168.1.200.
10. On the **Active/Active HA Binding** tab, for **Active/Active HA Binding**, select **primary** to bind the NAT rule to the firewall in active-primary state.
11. Click **OK**.

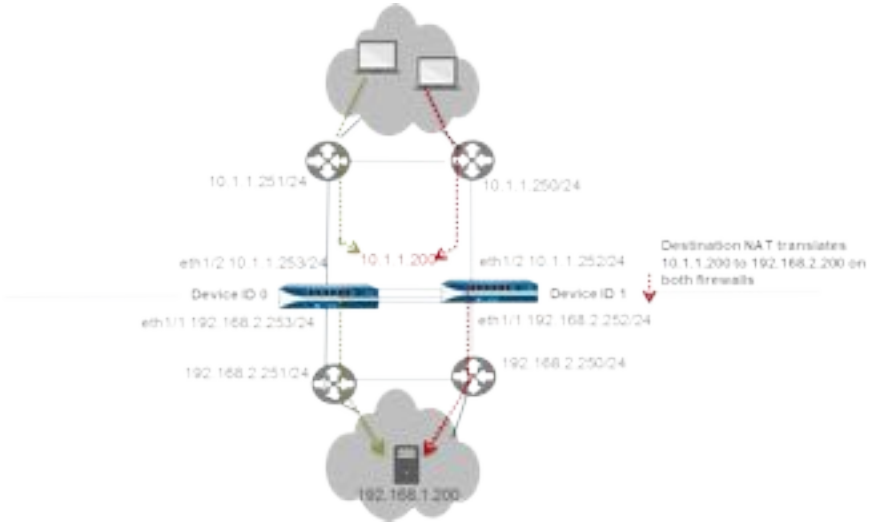
STEP 11 | **Commit** the configuration.

Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3

This Layer 3 interface example uses NAT in Active/Active HA Mode and ARP Load-Sharing. PA-3050-1 has Device ID 0 and its HA peer, PA-3050-2, has Device ID 1.

In this use case, both of the HA firewalls must respond to an ARP request for the destination NAT address. Traffic can arrive at either firewall from either WAN router in the untrust zone. Destination NAT translates the public-facing, shared IP address to the private IP address of the server. The configuration requires one destination NAT rule bound to both Device IDs so that both firewalls can respond to ARP requests.

© 2018 Palo Alto Networks, Inc.



STEP 1 | On PA-3050-2 (Device ID 1), perform Step 1 through Step 3 of Configure Active/Active HA.

STEP 2 | Enable active/active HA.

1. Select **Device** > **High Availability** > **General** > **Setup** and edit.
2. Select **Enable HA**.
3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63).
4. (Optional) Enter a **Description**.
5. For **Mode**, select **Active Active**.
6. Select **Device ID** to be **1**.
7. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default).
8. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall.
9. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall.
10. Click **OK**.

STEP 3 | Configure Active/Active HA.

Perform Step 6 through Step 15.

STEP 4 | Configure an HA virtual address.

1. Select **Device** > **High Availability** > **Active/Active Config** > **Virtual Address** and click **Add**.
2. Select **Interface** eth1/2.
3. Select **IPv4** and **Add** an **IPv4 Address** of 10.1.1.200.
4. For **Type**, select **ARP Load Sharing**, which configures the virtual IP address to be for both peers to use for ARP Load-Sharing.

STEP 5 | Configure ARP Load-Sharing.

The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.

1. For **Device Selection Algorithm**, select one of the following
   - **IP Modulo**—The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address.

© 2018 Palo Alto Networks, Inc.

- **IP Hash**—The firewall that will respond to ARP requests is based on a hash of the ARP requester's source IP address and destination IP address.

2. Click **OK**.

STEP 6 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

STEP 7 | Define HA Failover Conditions.

STEP 8 | **Commit** the configuration.

STEP 9 | Configure the peer firewall, PA-3050-1 (Device ID 0), with the same settings, except set the **Device ID** to **0** instead of **1**.

STEP 10 | Still on PA-3050-1 (Device ID 0), create the destination NAT rule for both Device ID 0 and Device ID 1.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a **Name** for the rule that in this example identifies it as a destination NAT rule for Layer 3 ARP.
3. For **NAT Type**, select **ipv4** (default).
4. On the **Original Packet**, for **Source Zone**, select **Any**.
5. For **Destination Zone**, select the Untrust zone you created for the external network.
6. Allow **Destination Interface**, **Service**, and **Source Address** to remain set to **Any**.
7. For **Destination Address**, specify 10.1.1.200.
8. For the **Translated Packet**, Source Address Translation remains None.
9. For **Destination Address Translation**, enter the private IP address of the destination server, in this example 192.168.1.200.
10. On the **Active/Active HA Binding** tab, for **Active/Active HA Binding**, select **both** to bind the NAT rule to both Device ID 0 and Device ID 1.
11. Click **OK**.

STEP 11 | **Commit** the configuration.

© 2018 Palo Alto Networks, Inc.

HA Firewall States

An HA firewall can be in one of the following states:

HA Firewall State	Occurs In	Description
Initial	A/P or A/A	Transient state of a firewall when it joins the HA pair. The firewall remains in this state after boot-up until it discovers a peer and negotiations begins. After a timeout, the firewall becomes active if HA negotiation has not started.
Active	A/P	State of the active firewall in an active/passive configuration.
Passive	A/P	State of the passive firewall in an active/passive configuration. The passive firewall is ready to become the active firewall with no disruption to the network. Although the passive firewall is not processing other traffic: • If passive link state auto is configured, the passive firewall is running routing protocols, monitoring link and path state, and the passive firewall will pre-negotiate LACP and LLDP if LACP and LLDP pre-negotiation are configured, respectively. • The passive firewall is synchronizing flow state, runtime objects, and configuration. • The passive firewall is monitoring the status of the active firewall using the hello protocol.
Active-Primary	A/A	In an active/active configuration, state of the firewall that connects to User-ID agents, runs DHCP server and DHCP relay, and matches NAT and PBF rules with the Device ID of the active-primary firewall. A firewall in this state can own sessions and set up sessions.
Active-Secondary	A/A	In an active/active configuration, state of the firewall that connects to User-ID agents, runs DHCP server, and matches NAT and PBF rules with the Device ID of the active-secondary firewall. A firewall in active-secondary state does not support DHCP relay. A firewall in this state can own sessions and set up sessions.
Tentative	A/A	State of a firewall (in an active/active configuration) caused by one of the following: • Failure of a firewall. • Failure of a monitored object (a link or path). • The firewall leaves suspended or non-functional state. A firewall in tentative state synchronizes sessions and configurations from the peer. • In a virtual wire deployment, when a firewall enters tentative state due to a path failure and receives a packet to forward, it sends the packet to the peer firewall over the HA3 link for processing. The peer firewall processes the packet and sends it back over the HA3

© 2018 Palo Alto Networks, Inc.

HA Firewall State	Occurs In	Description
		link to the firewall to be sent out the egress interface. This behavior preserves the forwarding path in a virtual wire deployment. • In a Layer 3 deployment, when a firewall in tentative state receives a packet, it sends that packet over the HA3 link for the peer firewall to own or set up the session. Depending on the network topology, this firewall either sends the packet out to the destination or sends it back to the peer in tentative state for forwarding. After the failed path or link clears or as a failed firewall transitions from tentative state to active-secondary state, the **Tentative Hold Time** is triggered and routing convergence occurs. The firewall attempts to build routing adjacencies and populate its route table before processing any packets. Without this timer, the recovering firewall would enter active-secondary state immediately and would blackhole packets because it would not have the necessary routes. When a firewall leaves suspended state, it goes into tentative state for the **Tentative Hold Time** after links are up and able to process incoming packets. **Tentative Hold Time range (sec)** can be disabled (which is 0 seconds) or in the range 10-600; default is 60.
Non-functional	A/P or A/A	Error state due to a dataplane failure or a configuration mismatch, such as only one firewall configured for packet forwarding, VR sync or QoS sync. In active/passive mode, all of the causes listed for Tentative state cause non-functional state.
Suspended	A/P or A/A	Administratively disabled state. In this state, an HA firewall cannot participate in the HA election process.

© 2018 Palo Alto Networks, Inc.

Reference: HA Synchronization

If you have enabled configuration synchronization on both peers in an HA pair, most of the configuration settings you configure on one peer will automatically sync to the other peer upon commit. To avoid configuration conflicts, always make configuration changes on the active (active/passive) or active-primary (active/active) peer and wait for the changes to sync to the peer before making any additional configuration changes.

Only committed configurations synchronize between HA peers. Any configuration in the commit queue at the time of an HA sync will not be synchronized.

The following topics identify which configuration settings you must configure on each firewall independently (these settings are not synchronized from the HA peer).

- What Settings Don't Sync in Active/Passive HA? on page 297
- What Settings Don't Sync in Active/Active HA? on page 299
- Synchronization of System Runtime Information on page 302

What Settings Don't Sync in Active/Passive HA?

You must configure the following settings on each firewall in an HA pair in an active/passive deployment. These settings do not sync from one peer to another.

Configuration Item	What Doesn't Sync in Active/Passive?
Management Interface Settings	All management configuration settings must be configured individually on each firewall, including: • **Device** > **Setup** > **Management** > **General Settings**—Hostname, Domain, Login Banner, SSL/TLS Service Profile, Time Zone, Locale, Date, Time, Latitude, Longitude. *The configuration for the associated SSL/TLS Service profile (Device > Certificate Management > SSL/TLS Service Profile and the associated certificates (Device > Certificate Management > Certificates) is synchronized. It is just the setting of which SSL/TLS Service Profile to use on the Management interface that does not sync.* • **Device** > **Setup** > **Management** > **Management Interface Settings**—IP Type, IP Address, Netmask, Default Gateway, IPv6 Address/Prefix Length, Default IPv6 Gateway, Speed, MTU, and Services (HTTP, HTTP OCSP, HTTPS, Telnet, SSH, Ping, SNMP, User-ID, User-ID Syslog Listener-SSL, User-ID Syslog Listener-UDP)
Multi-vsys Capability	You must activate the Virtual Systems license on each firewall in the pair to increase the number of virtual systems beyond the base number provided by default on PA-3000 Series, PA-5000 Series, PA-5200 Series, and PA-7000 Series firewalls. You must also enable **Multi Virtual System Capability** on each firewall (**Device** > **Setup** > **Management** > **General Settings**).

© 2018 Palo Alto Networks, Inc.

Configuration Item	What Doesn't Sync in Active/Passive?
Administrator Authentication Settings	You must define the authentication profile and certificate profile for administrative access to the firewall locally on each firewall (**Device** > **Setup** > **Management** > **Authentication**).
Panorama Settings	Set the following Panorama settings on each firewall (**Device** > **Setup** > **Management** > **Panorama Settings**). • **Panorama Servers** • **Disable Panorama Policy and Objects** and **Disable Device and Network Template**
SNMP	**Device** > **Setup** > **Operations** > **SNMP Setup**
Services	**Device** > **Setup** > **Services**
Global Service Routes	**Device** > **Setup** > **Services** > **Service Route Configuration**
Telemetry and Threat Intelligence Settings	**Device** > **Setup** > **Telemetry and Threat Intelligence**
Data Protection	**Device** > **Setup** > **Content-ID** > **Manage Data Protection**
Jumbo Frames	**Device** > **Setup** > **Session** > **Session Settings** > **Enable Jumbo Frame**
Forward Proxy Server Certificate Settings	**Device** > **Setup** > **Session** > **Decryption Settings** > **SSL Forward Proxy Settings**
Master Key Secured by HSM	**Device** > **Setup** > **HSM** > **Hardware Security Module Provider** > **Master Key Secured by HSM**
Log Export Settings	**Device** > **Scheduled Log Export**
Software Updates	With software updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer (**Device** > **Software**).
GlobalProtect Agent Package	With GlobalProtect client updates, you can either download and install them separately on each firewall, or download them to one peer and sync the update to the other peer. You must activate separately on each peer (**Device** > **GlobalProtect Client**).
Content Updates	With content updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer (**Device** > **Dynamic Updates**).
Licenses/ Subscriptions	**Device** > **Licenses**
Support Subscription	**Device** > **Support**

© 2018 Palo Alto Networks, Inc.

Configuration Item	What Doesn't Sync in Active/Passive?
Master Key	The master key must be identical on each firewall in the HA pair, but you must manually enter it on each firewall (**Device** > **Master Key and Diagnostics**). Before changing the master key, you must disable config sync on both peers (**Device** > **High Availability** > **General** > **Setup** and clear the **Enable Config Sync** check box) and then re-enable it after you change the keys.
Reports, logs, and Dashboard Settings	Log data, reports, and Dashboard data and settings (column display, widgets) are not synced between peers. Report configuration settings, however, are synced.
HA settings	**Device** > **High Availability**

What Settings Don't Sync in Active/Active HA?

You must configure the following settings on each firewall in an HA pair in an active/active deployment. These settings do not sync from one peer to another.

Configuration Item	What Doesn't Sync in Active/Active?
Management Interface Settings	You must configure all management settings individually on each firewall, including: • **Device** > **Setup** > **Management** > **General Settings**—Hostname, Domain, Login Banner, SSL/TLS Service Profile, Time Zone, Locale, Date, Time, Latitude, Longitude.  *The configuration for the associated SSL/TLS Service profile (Device > Certificate Management > SSL/TLS Service Profile and the associated certificates (Device > Certificate Management > Certificates) is synchronized. It is just the setting of which SSL/TLS Service Profile to use on the Management interface that does not sync.* • **Device** > **Setup** > **Management** > **Management Interface Settings**—IP Address, Netmask, Default Gateway, IPv6 Address/Prefix Length, Default IPv6 Gateway, Speed, MTU, and Services (HTTP, HTTP OCSP, HTTPS, Telnet, SSH, Ping, SNMP, User-ID, User-ID Syslog Listener-SSL, User-ID Syslog Listener-UDP)
Multi-vsys Capability	You must activate the Virtual Systems license on each firewall in the pair to increase the number of virtual systems beyond the base number provided by default on PA-3000 Series, PA-5000 Series, PA-5200 Series, and PA-7000 Series firewalls. You must also enable **Multi Virtual System Capability** on each firewall (**Device** > **Setup** > **Management** > **General Settings**).

© 2018 Palo Alto Networks, Inc.

Configuration Item	What Doesn't Sync in Active/Active?
Administrator Authentication Settings	You must define the authentication profile and certificate profile for administrative access to the firewall locally on each firewall (**Device** > **Setup** > **Management** > **Authentication**).
Panorama Settings	Set the following Panorama settings on each firewall (**Device** > **Setup** > **Management** > **Panorama Settings**). • **Panorama Servers** • **Disable Panorama Policy and Objects** and **Disable Device and Network Template**
SNMP	**Device** > **Setup** > **Operations** > **SNMP Setup**
Services	**Device** > **Setup** > **Services**
Global Service Routes	**Device** > **Setup** > **Services** > **Service Route Configuration**
Telemetry and Threat Intelligence Settings	**Device** > **Setup** > **Telemetry and Threat Intelligence**
Data Protection	**Device** > **Setup** > **Content-ID** > **Manage Data Protection**
Jumbo Frames	**Device** > **Setup** > **Session** > **Session Settings** > **Enable Jumbo Frame**
Forward Proxy Server Certificate Settings	**Device** > **Setup** > **Session** > **Decryption Settings** > **SSL Forward Proxy Settings**
HSM Configuration	**Device** > **Setup** > **HSM**
Log Export Settings	**Device** > **Scheduled Log Export**
Software Updates	With software updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer (**Device** > **Software**).
GlobalProtect Agent Package	With GlobalProtect client updates, you can either download and install them separately on each firewall, or download them to one peer and sync the update to the other peer. You must activate separately on each peer (**Device** > **GlobalProtect Client**).
Content Updates	With content updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer (**Device** > **Dynamic Updates**).
Licenses/Subscriptions	**Device** > **Licenses**
Support Subscription	**Device** > **Support**
Ethernet Interface IP Addresses	All Ethernet interface configuration settings sync except for the IP address (**Network** > **Interface** > **Ethernet**).

© 2018 Palo Alto Networks, Inc.

Configuration Item	What Doesn't Sync in Active/Active?
Loopback Interface IP Addresses	All Loopback interface configuration settings sync except for the IP address (**Network** > **Interface** > **Loopback**).
Tunnel Interface IP Addresses	All Tunnel interface configuration settings sync except for the IP address (**Network** > **Interface** > **Tunnel**).
LACP System Priority	Each peer must have a unique LACP System ID in an active/active deployment (**Network** > **Interface** > **Ethernet** > **Add Aggregate Group** > **System Priority**).
VLAN Interface IP Address	All VLAN interface configuration settings sync except for the IP address (**Network** > **Interface** > **VLAN**).
Virtual Routers	Virtual router configuration synchronizes only if you have enabled VR Sync (**Device** > **High Availability** > **Active/Active Config** > **Packet Forwarding**). Whether or not to do this depends on your network design, including whether you have asymmetric routing.
IPSec Tunnels	IPSec tunnel configuration synchronization is dependent on whether you have configured the Virtual Addresses to use Floating IP addresses (**Device** > **High Availability** > **Active/Active Config** > **Virtual Address**). If you have configured a floating IP address, these settings sync automatically. Otherwise, you must configure these settings independently on each peer.
GlobalProtect Portal Configuration	GlobalProtect portal configuration synchronization is dependent on whether you have configured the Virtual Addresses to use Floating IP addresses (**Network** > **GlobalProtect** > **Portals**). If you have configured a floating IP address, the GlobalProtect portal configuration settings sync automatically. Otherwise, you must configure the portal settings independently on each peer.
GlobalProtect Gateway Configuration	GlobalProtect gateway configuration synchronization is dependent on whether you have configured the Virtual Addresses to use Floating IP addresses (**Network** > **GlobalProtect** > **Gateways**). If you have configured a floating IP address, the GlobalProtect gateway configuration settings sync automatically. Otherwise, you must configure the gateway settings independently on each peer.
QoS	QoS configuration synchronizes only if you have enabled **QoS Sync** (**Device** > **High Availability** > **Active/Active Config** > **Packet Forwarding**). You might choose not to sync QoS setting if, for example, you have different bandwidth on each link or different latency through your service providers.
LLDP	No LLDP state or individual firewall data is synchronized in an active/active configuration (**Network** > **Network Profiles** > **LLDP**).
BFD	No BFD configuration or BFD session data is synchronized in an active/active configuration (**Network** > **Network Profiles** > **BFD Profile**).
IKE Gateways	IKE gateway configuration synchronization is dependent on whether you have configured the Virtual Addresses to use floating IP addresses (**Network** > **IKE Gateways**). If you have configured a floating IP address, the

© 2018 Palo Alto Networks, Inc.

Configuration Item	What Doesn't Sync in Active/Active?
	IKE gateway configuration settings sync automatically. Otherwise, you must configure the IKE gateway settings independently on each peer.
Master Key	The master key must be identical on each firewall in the HA pair, but you must manually enter it on each firewall (**Device** > **Master Key and Diagnostics**). Before changing the master key, you must disable config sync on both peers (**Device** > **High Availability** > **General** > **Setup** and clear the **Enable Config Sync** check box) and then re-enable it after you change the keys.
Reports, logs, and Dashboard Settings	Log data, reports, and dashboard data and settings (column display, widgets) are not synced between peers. Report configuration settings, however, are synced.
HA settings	• **Device** > **High Availability** • (The exception is **Device** > **High Availability** > **Active/Active Configuration** > **Virtual Addresses**, which do sync.)

Synchronization of System Runtime Information

The following table summarizes what system runtime information is synchronized between HA peers.

Runtime Information	Config Synced?		HA Link	Details
	A/P	A/A		
Management Plane				
User to Group Mappings	Yes	Yes	HA1	
DHCP Lease (as server)	Yes	Yes	HA1	
DNS Cache	No	No	N/A	
FQDN Refresh	No	No	N/A	
IKE Keys (phase 2)	Yes	Yes	HA1	
BrightCloud URL Database	No	No	N/A	
BrightCloud URL Cache	No	No	N/A	This feature is disabled by default and must be enabled separately on each HA peer.
BrightCloud Bloom Filter	No	No	N/A	This feature is disabled by default and must be enabled separately on each HA peer.
PAN-DB URL Cache	Yes	No	HA1	This is synchronized upon database backup to disk

© 2018 Palo Alto Networks, Inc.

Runtime Information	Config Synced?		HA Link	Details
	A/P	A/A		
				(every eight hours, when URL database version updates), or when the firewall reboots.
Content (manual sync)	Yes	Yes	HA1	
PPPoE, PPPoE Lease	Yes	Yes	HA1	
DHCP Client Settings and Lease	Yes	Yes	HA1	
SSL VPN Logged in User List	Yes	Yes	HA1	
Forward Information Base (FIB)	Yes	Yes	HA1	
Dataplane				
Session Table	Yes	Yes	HA2	• Active/passive peers do not sync ICMP or host session information. • Active/active peers do not sync host session, multicast session, or BFD session information.
ARP Table	Yes	No	HA2	
Neighbor Discovery (ND) Table	Yes	No	HA2	
MAC Table	Yes	No	HA2	
IPSec Sequence Number (anti-replay)	Yes	Yes	HA2	
DoS Block List Entries	No	No	N/A	
User to IP Address Mappings	Yes	Yes	HA2	
Virtual MAC	Yes	Yes	HA2	

© 2018 Palo Alto Networks, Inc.

Monitoring

To forestall potential issues and to accelerate incidence response when needed, the firewall provides intelligence about traffic and user patterns using customizable and informative reports. The dashboard, Application Command Center (ACC), reports, and logs on the firewall allow you to monitor activity on your network. You can monitor the logs and filter the information to generate reports with predefined or customized views. For example, you can use the predefined templates to generate reports on user activities or analyze the reports and logs to interpret unusual behavior on your network and generate a custom report on the traffic pattern. For a visually engaging presentation of network activity, the dashboard and the ACC include widgets, charts, and tables with which you can interact to find the information you care about. In addition, you can configure the firewall to forward monitored information as email notifications, syslog messages, SNMP traps, and NetFlow records to external services.

- Use the Dashboard on page 307
- Use the Application Command Center on page 309
- Use the App Scope Reports on page 328
- Use the Automated Correlation Engine on page 334
- Take Packet Captures on page 339
- Monitor Applications and Threats on page 350
- View and Manage Logs on page 351
- Monitor Block List on page 360
- View and Manage Reports on page 361
- Use External Services for Monitoring on page 377
- Configure Log Forwarding on page 378
- Configure Email Alerts on page 381
- Use Syslog for Monitoring on page 382
- SNMP Monitoring and Traps on page 418
- Forward Logs to an HTTP(S) Destination on page 433
- NetFlow Monitoring on page 436

© 2018 Palo Alto Networks, Inc.

Use the Dashboard

The **Dashboard** tab widgets show general firewall information, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available widgets are displayed by default, but each administrator can remove and add individual widgets, as needed. Click the refresh icon to update the dashboard or an individual widget. To change the automatic refresh interval, select an interval from the drop-down (**1 min**, **2 mins**, **5 mins**, or **Manual**). To add a widget to the dashboard, click the widget drop-down, select a category and then the widget name. To delete a widget, click in the title bar. The following table describes the dashboard widgets.

Dashboard Charts	Descriptions
Top Applications	Displays the applications with the most sessions. The block size indicates the relative number of sessions (mouse-over the block to view the number), and the color indicates the security risk—from green (lowest) to red (highest). Click an application to view its application profile.
Top High Risk Applications	Similar to Top Applications, except that it displays the highest-risk applications with the most sessions.
General Information	Displays the firewall name, model, PAN-OS software version, the application, threat, and URL filtering definition versions, the current date and time, and the length of time since the last restart.
Interface Status	Indicates whether each interface is up (green), down (red), or in an unknown state (gray).
Threat Logs	Displays the threat ID, application, and date and time for the last 10 entries in the Threat log. The threat ID is a malware description or URL that violates the URL filtering profile.
Config Logs	Displays the administrator username, client (Web or CLI), and date and time for the last 10 entries in the Configuration log.
Data Filtering Logs	Displays the description and date and time for the last 60 minutes in the Data Filtering log.
URL Filtering Logs	Displays the description and date and time for the last 60 minutes in the URL Filtering log.
System Logs	Displays the description and date and time for the last 10 entries in the System log. *A **`Config installed`** entry indicates configuration changes were committed successfully.*
System Resources	Displays the Management CPU usage, Data Plane usage, and the Session Count, which displays the number of sessions established through the firewall.

© 2018 Palo Alto Networks, Inc.

Dashboard Charts	Descriptions
Logged In Admins	Displays the source IP address, session type (Web or CLI), and session start time for each administrator who is currently logged in.
ACC Risk Factor	Displays the average risk factor (1 to 5) for the network traffic processed over the past week. Higher values indicate higher risk.
High Availability	If high availability (HA) is enabled, indicates the HA status of the local and peer firewall—green (active), yellow (passive), or black (other). For more information about HA, see High Availability on page 249.
Locks	Shows configuration locks taken by administrators.

© 2018 Palo Alto Networks, Inc.

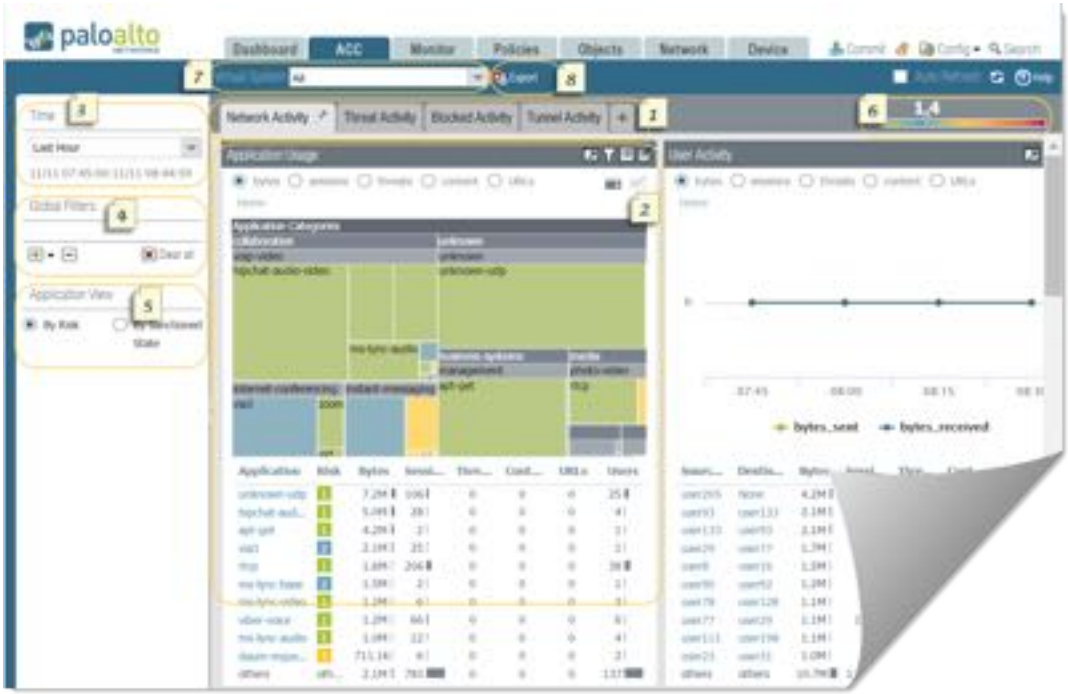
Use the Application Command Center

The Application Command Center (ACC) is an interactive, graphical summary of the applications, users, URLs, threats, and content traversing your network.The ACC uses the firewall logs to provide visibility into traffic patterns and actionable information on threats. The ACC layout includes a tabbed view of network activity, threat activity, and blocked activity and each tab includes pertinent widgets for better visualization of network traffic. The graphical representation allows you to interact with the data and visualize the relationships between events on the network, so that you can uncover anomalies or find ways to enhance your network security rules. For a personalized view of your network, you can also add a custom tab and include widgets that allow you to drill down into the information that is most important to you.

- ACC—First Look on page 309
- ACC Tabs on page 311
- ACC Widgets on page 312 (Widget Descriptions on page 314)
- ACC Filters on page 318
- Interact with the ACC on page 319
- Use Case: ACC—Path of Information Discovery on page 322

ACC—First Look

Take a quick tour of the ACC.



ACC—First Look		
1	**Tabs**	The ACC includes three predefined tabs that provide visibility into network traffic, threat activity, and blocked activity. For information on each tab, see ACC Tabs.

© 2018 Palo Alto Networks, Inc.

ACC—First Look		
2	**Widgets**	Each tab includes a default set of widgets that best represent the events/trends associated with the tab. The widgets allow you to survey the data using the following filters: • bytes (in and out) • sessions • content (files and data) • URL categories • threats (and count) For information on each widget, see ACC Widgets.
3	**Time**	The charts or graphs in each widget provide a summary and historic view. You can choose a custom range or use the predefined time periods that range from the last 15 minutes up to the last 30 days or last 30 calendar days. The selected time period applies across all tabs in the ACC. The time period used to render data, by default, is the **Last Hour** updated in 15 minute intervals. The date and time interval are displayed onscreen, for example at 11:40, the time range is 01/12 10:30:00-01/12 11:29:59.
4	**Global Filters**	The Global Filters allow you to set the filter across all widgets and all tabs. The charts/graphs apply the selected filters before rendering the data. For information on using the filters, see ACC Filters.
5	Application View	The application view allows you filter the ACC view by either the sanctioned and unsanctioned applications in use on your network, or by the risk level of the applications in use on your network. Green indicates sanctioned applications, blue unsanctioned applications, and yellow indicates applications that are partially sanctioned. Partially sanctioned applications are those that have a mixed sanctioned state; it indicates that the application is inconsistently tagged as sanctioned, for example it might be sanctioned on one or more virtual systems on a firewall enabled for multiple virtual systems or across one or more firewalls within a device group on Panorama.
6	**Risk Factor**	The risk factor (1=lowest to 5=highest) indicates the relative risk based on the applications used on your network. The risk factor uses a variety of factors to assess the associated risk levels, such as whether the application can share files, is it prone to misuse or does it try to evade firewalls, it also factors in the threat activity and malware as seen through the number

© 2018 Palo Alto Networks, Inc.

ACC—First Look		
		of blocked threats, compromised hosts or traffic to malware hosts/domains.
7	**Source**	The data used for the ACC display. The options vary on the firewall and on Panorama. On the firewall, if enabled for multiple virtual systems, you can use the **Virtual System** drop-down to change the ACC display to include data from all virtual systems or just a selected virtual system. On Panorama, you can select the **Device Group** drop-down to change the ACC display to include data from all device groups or just a selected device group. Additionally, on Panorama, you can change the **Data Source** as **Panorama** data or **Remote Device Data**. **Remote Device Data** is only available when all the managed firewalls are on PAN-OS 7.0.0 or later. When you filter the display for a specific device group, **Panorama** data is used as the data source**.**
8	**Export**	You can export the widgets displayed in the currently selected tab as a PDF. The PDF is downloaded and saved to the downloads folder associated with your web browser, on your computer.

ACC Tabs

The ACC includes the following predefined tabs for viewing network activity, threat activity, and blocked activity.

Tab	Description
Network Activity	Displays an overview of traffic and user activity on your network including: • Top applications in use • Top users who generate traffic (with a drill down into the bytes, content, threats or URLs accessed by the user) • Most used security rules against which traffic matches occur In addition, you can also view network activity by source or destination zone, region, or IP address, ingress or egress interfaces, and GlobalProtect host information such as the operating systems of the devices most commonly used on the network.
Threat Activity	Displays an overview of the threats on the network, focusing on the top threats: vulnerabilities, spyware, viruses, hosts visiting malicious domains or URLs, top WildFire submissions by file type and application, and applications that use non-standard ports. The Compromised Hosts widget in this tab (the widget is supported on some platforms only), supplements detection with better visualization techniques; it uses the information

© 2018 Palo Alto Networks, Inc.

Tab	Description
	from the correlated events tab (**Automated Correlation Engine** > **Correlated Events**) to present an aggregated view of compromised hosts on your network by source users/IP addresses and sorted by severity.
Blocked Activity	Focuses on traffic that was prevented from coming into the network. The widgets in this tab allow you to view activity denied by application name, username, threat name, blocked content—files and data that were blocked by a file blocking profile. It also lists the top security rules that were matched on to block threats, content, and URLs.
Tunnel Activity	Displays the activity of tunnel traffic that the firewall inspected based on your tunnel inspection policies. Information includes tunnel usage based on tunnel ID, monitor tag, user, and tunnel protocols such as Generic Routing Encapsulation (GRE), General Packet Radio Service (GPRS) Tunneling Protocol for User Data (GTP-U), and non-encrypted IPSec.

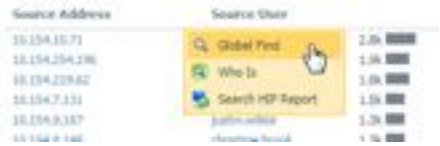
You can also Interact with the ACC on page 319 to create customized tabs with custom layout and widgets that meet your network monitoring needs, export the tab and share with another administrator.

ACC Widgets

The widgets on each tab are interactive; you can set the ACC Filters and drill down into the details for each table or graph, or customize the widgets included in the tab to focus on the information you need. For details on what each widget displays, see Widget Descriptions.



© 2018 Palo Alto Networks, Inc.

Widgets		
1	**View**	You can sort the data by bytes, sessions, threats, count, content, URLs, malicious, benign, files, applications, data, profiles, objects, users. The available options vary by widget.
2	**Graph**	The graphical display options are treemap, line graph, horizontal bar graph, stacked area graph, stacked bar graph, and map. The available options vary by widget; the interaction experience also varies with each graph type. For example, the widget for Applications using Non-Standard Ports allows you to choose between a treemap and a line graph. To drill down into the display, click into the graph. The area you click into becomes a filter and allows you to zoom into the selection and view more granular information on the selection.
3	**Table**	The detailed view of the data used to render the graph is provided in a table below the graph. You can interact with the table in several ways: • Click and set a local filter for an attribute in the table. The graph is updated and the table is sorted using the local filter. The information displayed in the graph and the table are always synchronized. • Hover over the attribute in the table and use the options available in the drop-down. 
4	**Actions**	**Maximize view**— Allows you enlarge the widget and view the table in a larger screen space and with more viewable information. **Set up local filters**—Allows you to add ACC Filters to refine the display within the widget. Use these filters to customize the widgets; these customizations are retained between logins. **Jump to logs**—Allows you to directly navigate to the logs (**Monitor** > **Logs** > **<log-type>** tab). The logs are filtered using the time period for which the graph is rendered. If you have set local and global filters, the log query concatenates the time period and the filters and only displays logs that match the combined filter set. **Export**—Allows you to export the graph as a PDF. The PDF is downloaded and saved on your computer. It is

© 2018 Palo Alto Networks, Inc.

Widgets		
		saved in the Downloads folder associated with your web browser.

Widget Descriptions

Each tab on the ACC includes a different set of widgets.

Widget	Description
Network Activity—Displays an overview of traffic and user activity on your network.	
Application Usage	The table displays the top ten applications used on your network, all the remaining applications used on the network are aggregated and displayed as other. The graph displays all applications by application category, sub category, and application. Use this widget to scan for applications being used on the network, it informs you about the predominant applications using bandwidth, session count, file transfers, triggering the most threats, and accessing URLs. Sort attributes: bytes, sessions, threats, content, URLs Charts available: treemap, area, column, line (the charts vary by the sort by attribute selected)
User Activity	Displays the top ten most active users on the network who have generated the largest volume of traffic and consumed network resources to obtain content. Use this widget to monitor top users on usage sorted on bytes, sessions, threats, content (files and patterns), and URLs visited. Sort attributes: bytes, sessions, threats, content, URLs Charts available: area, column, line (the charts vary by the sort by attribute selected)
Source IP Activity	Displays the top ten IP addresses or hostnames of the devices that have initiated activity on the network. All other devices are aggregated and displayed as other. Sort attributes: bytes, sessions, threats, content, URLs Charts available: area, column, line (the charts vary by the sort by attribute selected)
Destination IP Activity	Displays the IP addresses or hostnames of the top ten destinations that were accessed by users on the network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: area, column, line (the charts vary by the sort by attribute selected)
Source Regions	Displays the top ten regions (built-in or custom defined regions) around the world from where users initiated activity on your network. Sort attributes: bytes, sessions, threats, content, URLs

© 2018 Palo Alto Networks, Inc.

Widget	Description
	Charts available: map, bar
Destination Regions	Displays the top ten destination regions (built-in or custom defined regions) on the world map from where content is being accessed by users on the network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: map, bar
GlobalProtect Host Information	Displays information on the state of the hosts on which the GlobalProtect agent is running; the host system is a GlobalProtect client. This information is sourced from entries in the HIP match log that are generated when the data submitted by the GlobalProtect agent matches a HIP object or a HIP profile you have defined on the firewall. If you do not have HIP Match logs, this widget is blank. To learn how to create HIP objects and HIP profiles and use them as policy match criteria, see Configure HIP-Based Policy Enforcement. Sort attributes: profiles, objects, operating systems Charts available: bar
Rule Usage	Displays the top ten rules that have allowed the most traffic on the network. Use this widget to view the most commonly used rules, monitor the usage patterns, and to assess whether the rules are effective in securing your network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: line
Ingress Interfaces	Displays the firewall interfaces that are most used for allowing traffic into the network. Sort attributes: bytes, bytes sent, bytes received Charts available: line
Egress Interfaces	Displays the firewall interfaces that are most used by traffic exiting the network. Sort attributes: bytes, bytes sent, bytes received Charts available: line
Source Zones	Displays the zones that are most used for allowing traffic into the network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: line
Destination Zones	Displays the zones that are most used by traffic going outside the network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: line
Threat Activity—Displays an overview of the threats on the network	

© 2018 Palo Alto Networks, Inc.

Widget	Description
Compromised Hosts	Displays the hosts that are likely compromised on your network. This widget summarizes the events from the correlation logs. For each source user/IP address, it includes the correlation object that triggered the match and the match count, which is aggregated from the match evidence collated in the correlated events logs. For details see Use the Automated Correlation Engine on page 334. Available on the PA-3000 Series, PA-5000 Series, PA-5200 Series, PA-7000 Series, and Panorama. Sort attributes: severity (by default)
Hosts Visiting Malicious URLs	Displays the frequency with which hosts (IP address/hostnames) on your network have accessed malicious URLs. These URLs are known to be malware based on categorization in PAN-DB. Sort attributes: count Charts available: line
Hosts Resolving Malicious Domains	Displays the top hosts matching DNS signatures; hosts on the network that are attempting to resolve the hostname or domain of a malicious URL. This information is gathered from an analysis of the DNS activity on your network. It utilizes passive DNS monitoring, DNS traffic generated on the network, activity seen in the sandbox if you have configured DNS sinkhole on the firewall, and DNS reports on malicious DNS sources that are available to Palo Alto Networks customers. Sort attributes: count Charts available: line
Threat Activity	Displays the threats seen on your network. This information is based on signature matches in Antivirus, Anti-Spyware, and Vulnerability Protection profiles and viruses reported by WildFire. Sort attributes: threats Charts available: bar, area, column
WildFire Activity by Application	Displays the applications that generated the most WildFire submissions. This widget uses the malicious and benign verdict from the WildFire Submissions log. Sort attributes: malicious, benign Charts available: bar, line
WildFire Activity by File Type	Displays the threat vector by file type. This widget displays the file types that generated the most WildFire submissions and uses the malicious and benign verdict from the WildFire Submissions log. If this data is unavailable, the widget is empty. Sort attributes: malicious, benign Charts available: bar, line

© 2018 Palo Alto Networks, Inc.

Widget	Description
Applications using Non Standard Ports	Displays the applications that are entering your network on non-standard ports. If you have migrated your firewall rules from a port-based firewall, use this information to craft policy rules that allow traffic only on the default port for the application. Where needed, make an exception to allow traffic on a non-standard port or create a custom application. Sort attributes: bytes, sessions, threats, content, URLs Charts available: treemap, line
Rules Allowing Applications On Non Standard Ports	Displays the security policy rules that allow applications on non-default ports. The graph displays all the rules, while the table displays the top ten rules and aggregates the data from the remaining rules as other. This information helps you identify gaps in network security by allowing you to assess whether an application is hopping ports or sneaking into your network. For example, you can validate whether you have a rule that allows traffic on any port except the default port for the application. Say for example, you have a rule that allow DNS traffic on its *application-default* port (port 53 is the standard port for DNS). This widget will display any rule that allows DNS traffic into your network on any port except port 53. Sort attributes: bytes, sessions, threats, content, URLs Charts available: treemap, line
Blocked Activity—Focuses on traffic that was prevented from coming into the network	
Blocked Application Activity	Displays the applications that were denied on your network, and allows you to view the threats, content, and URLs that you kept out of your network. Sort attributes: threats, content, URLs Charts available: treemap, area, column
Blocked User Activity	Displays user requests that were blocked by a match on an Antivirus, Anti-spyware, File Blocking or URL Filtering profile attached to Security policy rule. Sort attributes: threats, content, URLs Charts available: bar, area, column
Blocked Threats	Displays the threats that were successfully denied on your network. These threats were matched on antivirus signatures, vulnerability signatures, and DNS signatures available through the dynamic content updates on the firewall. Sort attributes: threats Charts available: bar, area, column
Blocked Content	Displays the files and data that was blocked from entering the network. The content was blocked because security policy denied access based on criteria defined in a File Blocking security profile or a Data Filtering security profile. Sort attributes: files, data Charts available: bar, area, column

© 2018 Palo Alto Networks, Inc.

Widget	Description
Security Policies Blocking Activity	Displays the security policy rules that blocked or restricted traffic into your network. Because this widget displays the threats, content, and URLs that were denied access into your network, you can use it to assess the effectiveness of your policy rules. This widget does not display traffic that blocked because of deny rules that you have defined in policy. Sort attributes: threats, content, URLs Charts available: bar, area, column

ACC Filters

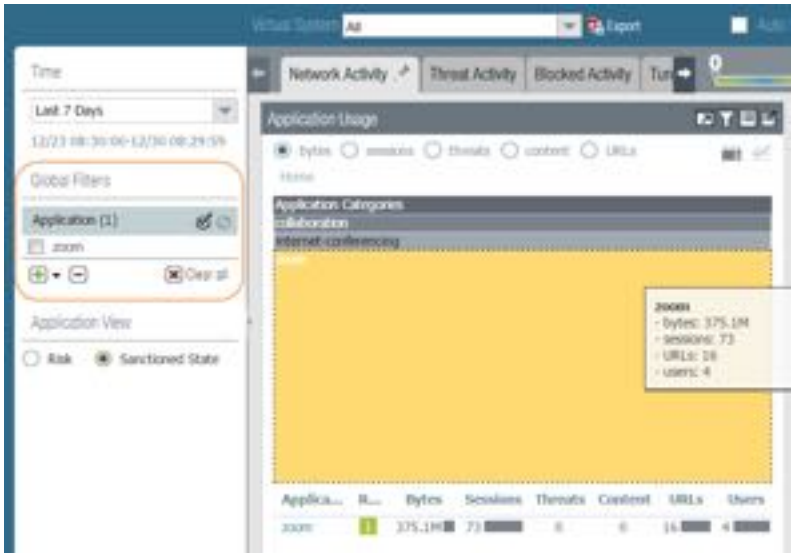
The graphs and tables on the ACC widgets allow you to use filters to narrow the scope of data that is displayed, so that you can isolate specific attributes and analyze information you want to view in greater detail. The ACC supports the simultaneous use of widget and global filters.

- **Widget Filters**—Apply a widget filter, which is a filter that is *local* to a specific widget. A widget filter allows you to interact with the graph and customize the display so that you can drill down in to the details and access the information you want to monitor on a specific widget. To create a widget filter that is persistent across reboots, you must use the **Set Local Filter** option.



- **Global filters**—Apply global filters across all the tabs in the ACC. A global filter allows you to pivot the display around the details you care about right now and exclude the unrelated information from the current display. For example, to view all events relating to a specific user and application, you can apply the username and the application as a global filter and view only information pertaining to the user and the application through all the tabs and widgets on the ACC. Global filters are not persistent.

© 2018 Palo Alto Networks, Inc.



You can apply global filters in three ways:

- **Set a global filter from a table**—Select an attribute from a table in any widget and apply the attribute as a global filter.
- **Add a widget filter to a global filter**—Hover over the attribute and click the arrow icon to the right of the attribute. This option allows you to elevate a local filter used in a widget, and apply the attribute globally to update the display across all the tabs on the ACC.
- **Define a global filter**—Define a filter using the **Global Filters** pane on the ACC.

See Interact with the ACC on page 319 for details on using these filters.

Interact with the ACC

To customize and refine the ACC display, you can add, delete, export and import tabs, add and delete widgets, set local and global filters, and interact with the widgets.

- Add a tab.
  1. Select the + icon along the list of tabs.
  2. Add a **View Name**. This name will be used as the name for the tab. You can add up to five tabs.

- Edit a tab.

  Select the tab, and click the pencil icon next to the tab name, to edit the tab. For example  .

  Editing a tab allows you to add or delete or reset the widgets that are displayed in the tab. You can also change the widget layout in the tab.

  To save the tab as the default tab, select .

- Export and Import tabs.
  1. Select the tab, and click the pencil icon next to the tab name, to edit the tab.
  2. Select the icon to export the current tab as a .txt file. You can share this .txt file with another administrator.

© 2018 Palo Alto Networks, Inc.

3. To import the tab as a new tab on another firewall, select the + icon along the list of tabs, and add a name and click the import icon, browse to select the .txt file.



- See what widgets are included in a tab.
  1. Select the tab, and click on the pencil icon to edit it.
  2. Select the **Add Widget** drop-down and verify the widgets that have the check boxes selected.

- Add a widget or a widget group.
  1. Add a new tab or edit a predefined tab.
  2. Select **Add Widget**, and then select the check box that corresponds to the widget you want to add. You can select up to a maximum of 12 widgets.
  3. (Optional) To create a 2-column layout, select **Add Widget Group**. You can drag and drop widgets into the 2-column display. As you drag the widget into the layout, a placeholder will display for you to drop the widget.

     You cannot name a widget group.

- Delete a tab or a widget group/ widget.
  1. To delete a custom tab, select the tab and click the X icon. 

     You cannot delete a predefined tab.
  2. To delete a widget group/widget, edit the tab and in the workspace section, click the [X] icon on the right. You cannot undo a deletion.

- Reset the default widgets in a tab.

  On a predefined tab, such as the **Blocked Activity** tab, you can delete one or more widgets. If you want to reset the layout to include the default set of widgets for the tab, edit the tab and click **Reset View**.

- Zoom in on the details in an area, column, or line graph.

  Watch how the zoom-in capability works.

  Click and drag an area in the graph to zoom in. For example, when you zoom into a line graph, it triggers a re-query and the firewall fetches the data for the selected time period. It is not a mere magnification.

- Use the table drop-down to find more information on an attribute.
  1. Hover over an attribute in a table to see the drop-down.
  2. Click into the drop-down to view the available options.
     - **Global Find**—Use Global Find to Search the Firewall or Panorama Management Server on page 66 for references to the attribute (username/IP address, object name, policy rule name, threat ID, or application name) anywhere in the candidate configuration.
     - **Value**—Displays the details of the threat ID, or application name, or address object.

© 2018 Palo Alto Networks, Inc.

- **Who Is**—Performs a domain name (*WHOIS*) lookup for the IP address. The lookup queries databases that store the registered users or assignees of an Internet resource.
- **Search HIP Report**—Uses the username or IP address to find matches in a HIP Match report.

- Set a widget filter.

  You can also click an attribute in the table (below the graph) to apply it as a widget filter.

  1. Select a widget and click the icon.
  2. Click the icon to add the filters you want to apply.
  3. Click **Apply**. These filters are persistent across reboots.

  The active widget filters are indicated next to the widget name.

- Negate a widget filter
  1. Click the icon to display the Setup Local Filters dialog.
  2. Add a filter, and then click the negate icon.

- Set a global filter from a table.

  Hover over an attribute in the table below the chart and click the arrow icon to the right of the attribute.



- Set a global filter using the Global Filters pane.

  Watch global filters in action.

  1. Locate the **Global Filters** pane on the left side of the ACC.



  2. Click the icon to view the list of filters you can apply.

- Promote a widget filter to a global filter.
  1. On any table in a widget, click the link for an attribute. This sets the attribute as a widget filter.
  2. To promote the filter to be a global filter, select the arrow to the right of the filter.

© 2018 Palo Alto Networks, Inc.



- Remove a filter.

  Click the icon to remove a filter.

  - For global filters: It is located in the Global Filters pane.
  - For widget filters: Click the icon to display the Setup Local Filters dialog, then select the filter, and click the icon.

- Clear all filters.

  - For global filters: Click the **Clear All** button under Global Filters.
  - For widget filters: Select a widget and click the icon. Then click the **Clear All** button in the Setup Local Filters dialog.

- See what filters are in use.

  - For global filters: The number of global filters applied are displayed on the left pane under Global Filters.
  - For widget filters: The number of widget filters applied on a widget are displayed next to the widget name. To view the filters, click the icon.

- Reset the display on a widget.

  - If you set a widget filter or drill into a graph, click the **Home** link to reset the display in the widget.



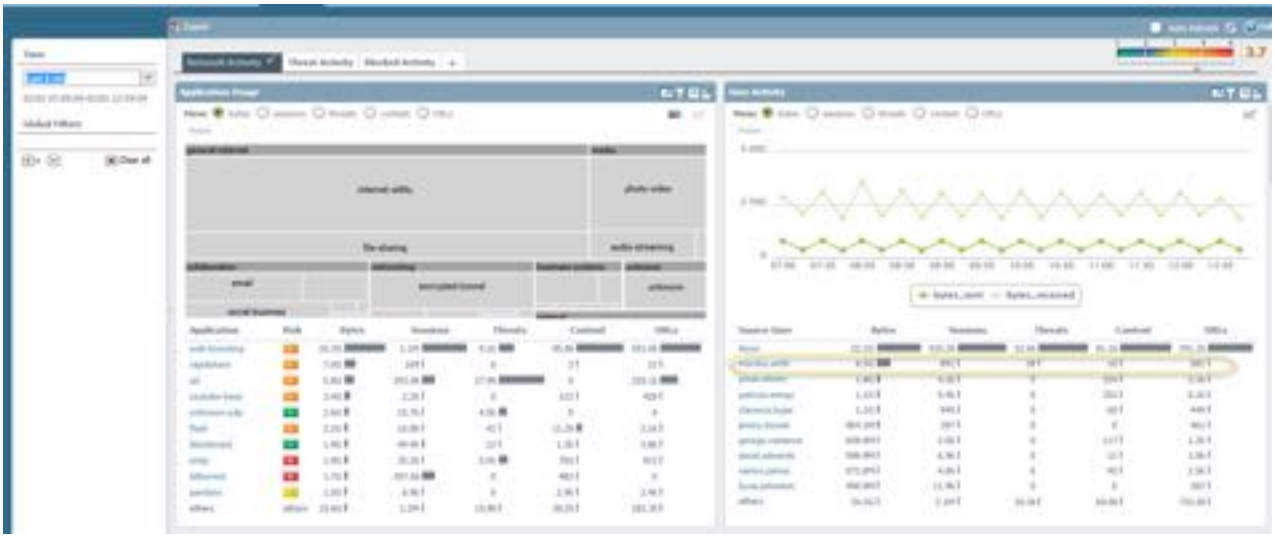
Use Case: ACC—Path of Information Discovery

The ACC has a wealth of information that you can use as a starting point for analyzing network traffic. Let's look at an example on using the ACC to uncover events of interest. This example illustrates how you can use the ACC to ensure that legitimate users can be held accountable for their actions, detect and track unauthorized activity, and detect and diagnose compromised hosts and vulnerable systems on your network.

The widgets and filters in the ACC give you the capability to analyze the data and filter the views based on events of interest or concern. You can trace events that pique your interest, directly export a PDF of a tab, access the raw logs, and save a personalized view of the activity that you want to track. These capabilities make it possible for you to monitor activity and develop policies and countermeasures for fortifying your network against malicious activity. In this section, you will Interact with the ACC on page 319 widgets across different tabs, drill down using widget filters, and pivot the ACC views using global filters, and export a PDF for sharing with incidence response or IT teams.

© 2018 Palo Alto Networks, Inc.

At first glance, you see the Application Usage and User Activity widgets in the **ACC** > **Network Activity** tab. The User Activity widget shows that user Marsha Wirth has transferred 718 Megabytes of data during the last hour. This volume is nearly six times more than any other user on the network. To see the trend over the past few hours, expand the **Time** period to the **Last 6 Hrs**, and now Marsha's activity has been 6.5 Gigabytes over 891 sessions and has triggered 38 threats signatures.



Because Marsha has transferred a large volume of data, apply her username as a global filter (ACC Filters on page 318) and pivot all the views in the ACC to Marsha's traffic activity.



The Application Usage tab now shows that the top application that Martha used was rapidshare, a Swiss-owned file-hosting site that belongs to the file-sharing URL category. For further investigation, add rapidshare as a global filter, and view Marsha's activity in the context of rapidshare.

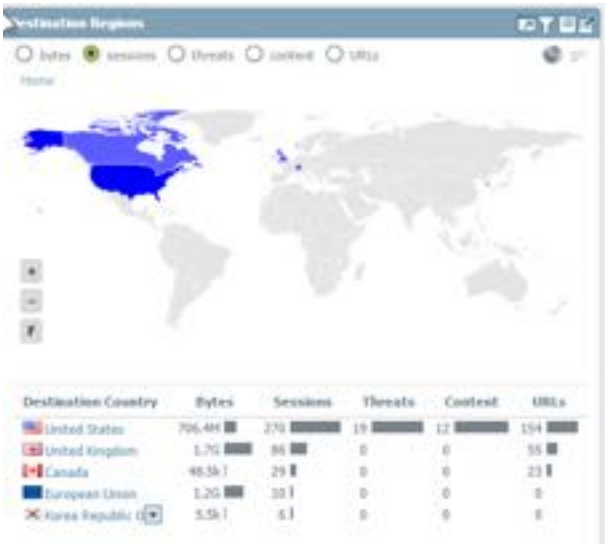
Consider whether you want to sanction rapidshare for company use. Should you allow uploads to this site and do you need a QoS policy to limit bandwidth?

To view which IP addresses Marsha has communicated with, check the **Destination IP Activity** widget, and view the data by bytes and by URLs.

© 2018 Palo Alto Networks, Inc.



To find out which countries Marsha communicated with, sort on **sessions** in the **Destination Regions** widget.



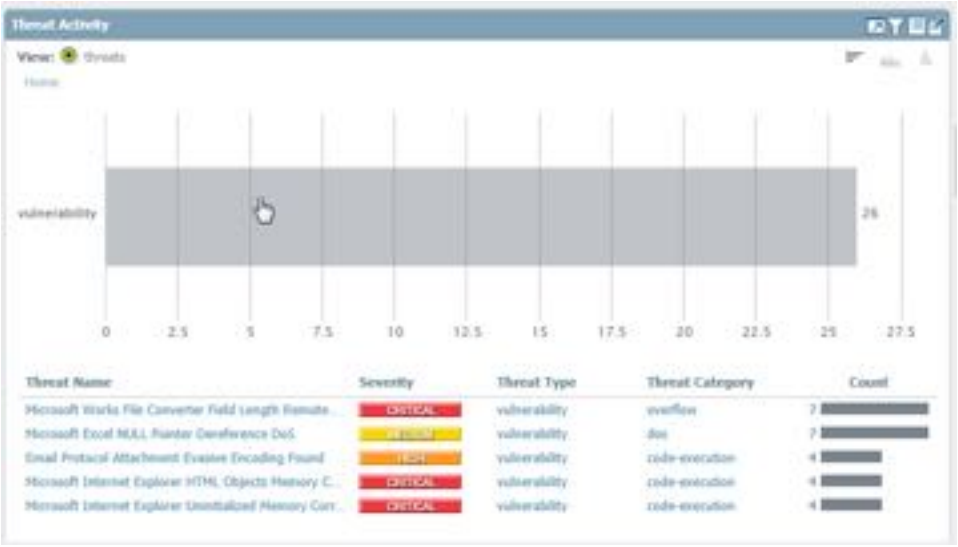
From this data, you can confirm that Marsha, a user on your network, has established sessions in Korea and the European Union, and she logged 19 threats in her sessions within the United States.

To look at Marsha's activity from a threat perspective, remove the global filter for rapidshare.



In the **Threat Activity** widget on the **Threat Activity** tab, view the threats. The widget displays that her activity had triggered a match for 26 vulnerabilities in the overflow, DoS and code-execution threat category. Several of these vulnerabilities are of critical severity.

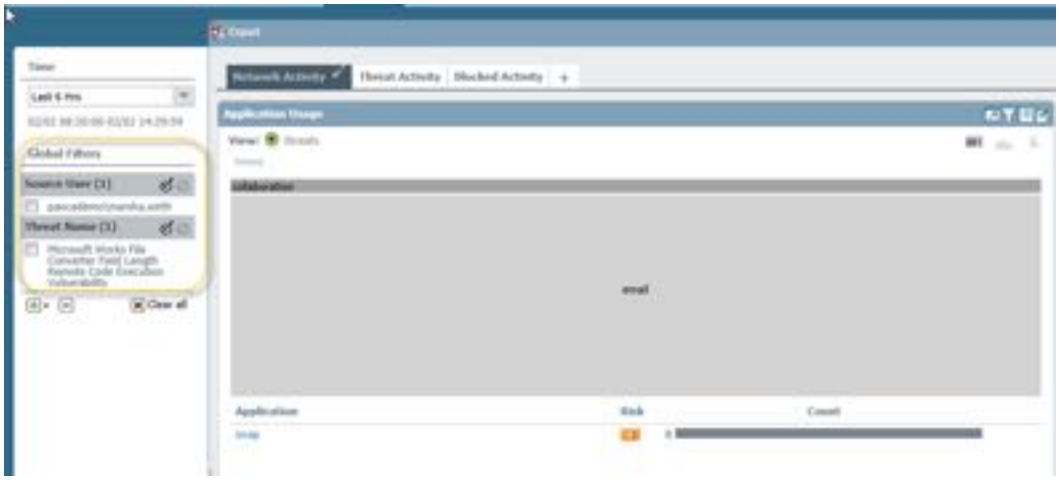
© 2018 Palo Alto Networks, Inc.



To further drill-down into each vulnerability, click into the graph and narrow the scope of your investigation. Each click automatically applies a local filter on the widget.



To investigate each threat by name, you can create a global filter for say, **Microsoft Works File Converter Field Length Remote Code Execution Vulnerability**. Then, view the **User Activity widget** in the **Network Activity** tab. The tab is automatically filtered to display threat activity for Marsha (notice the global filters in the screenshot).



© 2018 Palo Alto Networks, Inc.

Notice that this Microsoft code-execution vulnerability was triggered over email, by the imap application. You can now establish that Martha has IE vulnerabilities and email attachment vulnerabilities, and perhaps her computer needs to be patched. You can now either navigate to the **Blocked Threats** widget in the **Blocked Activity** tab to check how many of these vulnerabilities were blocked.

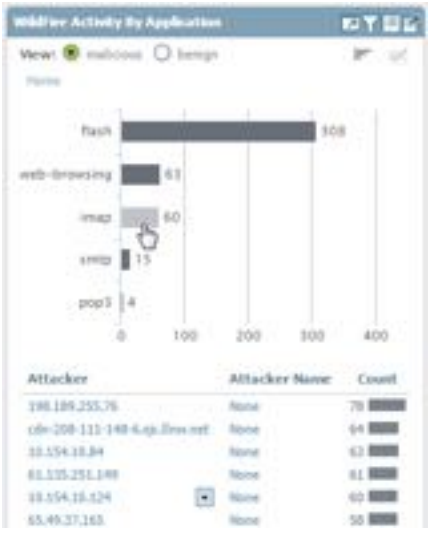
Or, you can check the **Rule Usage** widget on the **Network Activity** tab to discover how many vulnerabilities made it into your network and which security rule allowed this traffic, and navigate directly to the security rule using the **Global Find** capability.



Then, drill into why imap used a non-standard port 43206 instead of port 143, which is the default port for the application. Consider modifying the security policy rule to allow applications to only use the default port for the application, or assess whether this port should be an exception on your network.



To review if any threats were logged over imap, check Marsha's activity in the **WildFire Activity by Application** widget in the **Threat Activity** tab. You can confirm that Marsha had no malicious activity, but to verify that other no other user was compromised by the imap application, negate Marsha as a global filter and look for other users who triggered threats over imap.



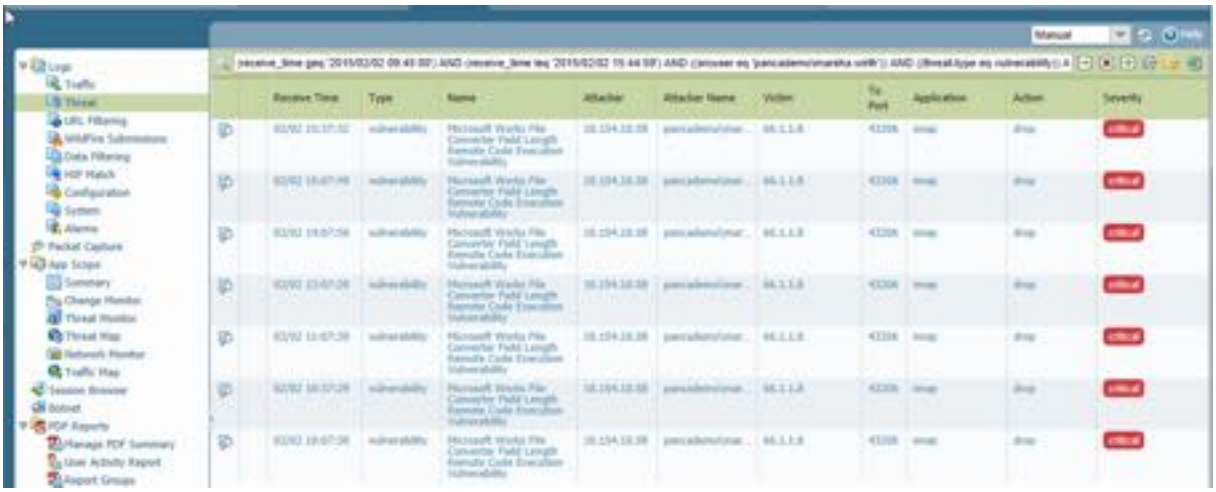
© 2018 Palo Alto Networks, Inc.

Click into the bar for imap in the graph and drill into the inbound threats associated with the application. To find out who an IP address is registered to, hover over the attacker IP address and select the **Who Is** link in the drop-down.



Because the session count from this IP address is high, check the **Blocked Content** and **Blocked Threats** widgets in the **Blocked Activity** tab for events related to this IP address. The **Blocked Activity** tab allows you to validate whether or not your policy rules are effective in blocking content or threats when a host on your network is compromised.

Use the **Export PDF** capability on the ACC to export the current view (create a snapshot of the data) and send it to an incidence response team. To view the threat logs directly from the widget, you can also click the icon to jump to the logs; the query is generated automatically and only the relevant logs are displayed onscreen (for example in **Monitor** > **Logs** > **Threat Logs**).



You have now used the ACC to review network data/trends to find which applications or users are generating the most traffic, and how many application are responsible for the threats seen on the network. You were able to identify which application(s), user(s) generated the traffic, determine whether the application was on the default port, and which policy rule(s) allowed the traffic into the network, and determine whether the threat is spreading laterally on the network. You also identified the destination IP addresses, geo-locations with which hosts on the network are communicating with. Use the conclusions from your investigation to craft goal-oriented policies that can secure users and your network.

© 2018 Palo Alto Networks, Inc.

Use the App Scope Reports

The App Scope reports provide visibility and analysis tools to help pinpoint problematic behavior, helping you understand changes in application usage and user activity, users and applications that take up most of the network bandwidth, and identify network threats.

With the App Scope reports, you can quickly see if any behavior is unusual or unexpected. Each report provides a dynamic, user-customizable window into the network; hovering the mouse over and clicking either the lines or bars on the charts opens detailed information about the specific application, application category, user, or source on the ACC. The App Scope charts on **Monitor** > **App Scope** give you the ability to:

- Toggle the attributes in the legend to only view chart details that you want to review. The ability to include or exclude data from the chart allows you to change the scale and review details more closely.
- Click into an attribute in a bar chart and drill down to the related sessions in the ACC. Click into an Application name, Application Category, Threat Name, Threat Category, Source IP address or Destination IP address on any bar chart to filter on the attribute and view the related sessions in the ACC.
- Export a chart or map to PDF or as an image. For portability and offline viewing, you can Export charts and maps as PDFs or PNG images.

The following App Scope reports are available:

- Summary Report on page 328
- Change Monitor Report on page 329
- Threat Monitor Report on page 330
- Threat Map Report on page 331
- Network Monitor Report on page 331
- Traffic Map Report on page 332

Summary Report

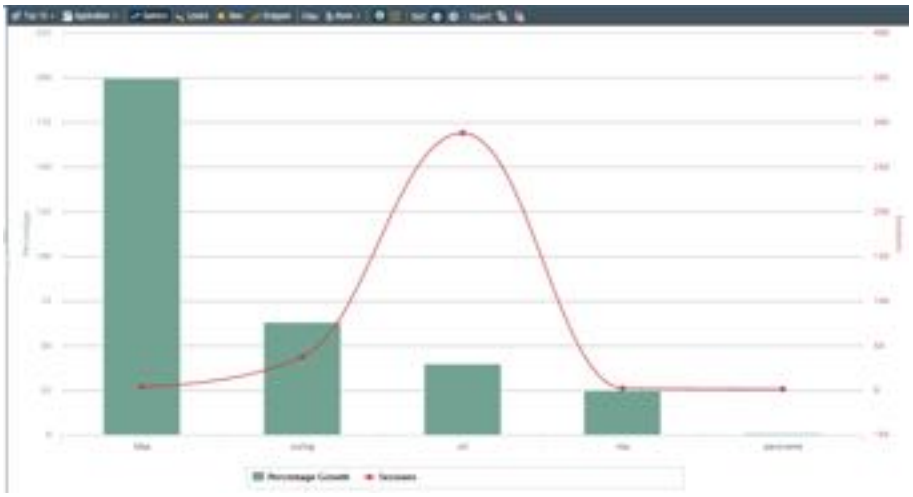
The App Scope Summary report (**Monitor** > **App Scope** > **Summary**) displays charts for the top five gainers, losers, and bandwidth consuming applications, application categories, users, and sources.



© 2018 Palo Alto Networks, Inc.

Change Monitor Report

The App Scope Change Monitor report (**Monitor** > **App Scope** > **Change Monitor**) displays changes over a specified time period. For example, the following chart displays the top applications that gained in use over the last hour as compared with the last 24-hour period. The top applications are determined by session count and sorted by percent.



The Change Monitor Report contains the following buttons and options.

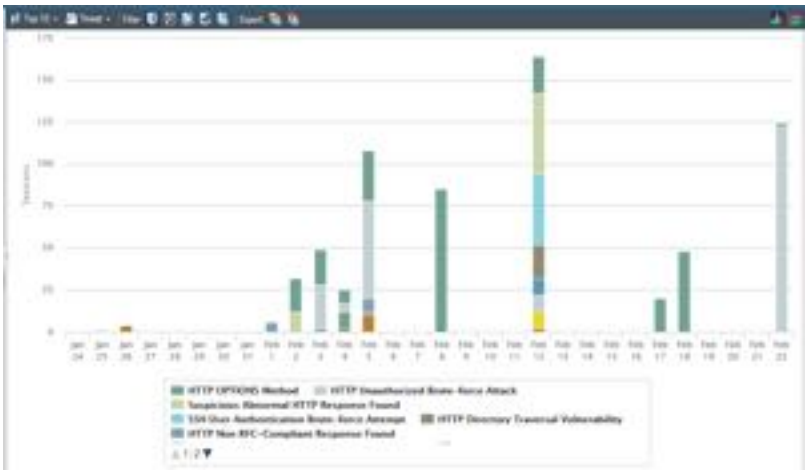
Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Application	Determines the type of item reported: Application, Application Category, Source, or Destination.
Gainers	Displays measurements of items that have increased over the measured period.
Losers	Displays measurements of items that have decreased over the measured period.
New	Displays measurements of items that were added over the measured period.
Dropped	Displays measurements of items that were discontinued over the measured period.
Filter	Applies a filter to display only the selected item. None displays all entries.
	Determines whether to display session or byte information.
Sort	Determines whether to sort entries by percentage or raw growth.

© 2018 Palo Alto Networks, Inc.

Button	Description
Export	Exports the graph as a .png image or as a PDF.
Compare	Specifies the period over which the change measurements are taken.

Threat Monitor Report

The App Scope Threat Monitor report (**Monitor** > **App Scope** > **Threat Monitor**) displays a count of the top threats over the selected time period. For example, the following figure shows the top 10 threat types over the last 6 hours.



Each threat type is color-coded as indicated in the legend below the chart. The Threat Monitor report contains the following buttons and options.

Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Threats	Determines the type of item measured: Threat, Threat Category, Source, or Destination.
Filter	Applies a filter to display only the selected type of items.
	Determines whether the information is presented in a stacked column chart or a stacked area chart.
Export	Exports the graph as a .png image or as a PDF.
	Specifies the period over which the measurements are taken.

© 2018 Palo Alto Networks, Inc.

Threat Map Report

The App Scope Threat Map report (**Monitor** > **App Scope** > **Threat Map**) shows a geographical view of threats, including severity. Each threat type is color-coded as indicated in the legend below the chart.

The firewall uses geolocation for creating threat maps. The firewall is placed at the bottom of the threat map screen, if you have not specified the geolocation coordinates (**Device** > **Setup** > **Management**, General Settings section) on the firewall.



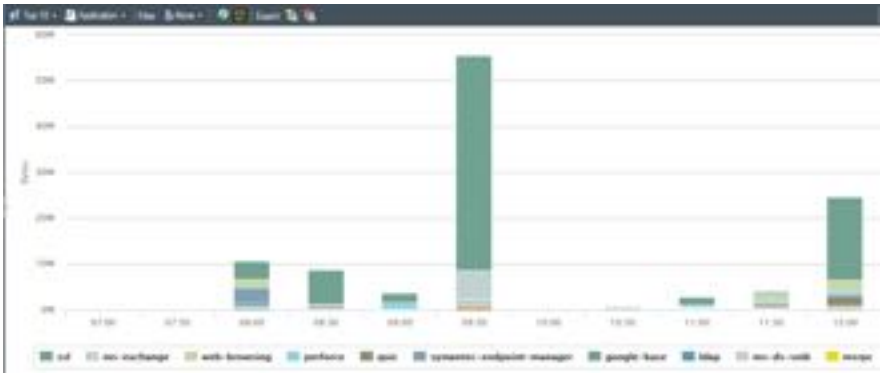
The Threat Map report contains the following buttons and options.

Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Incoming threats	Displays incoming threats.
Outdoing threats	Displays outgoing threats.
Filer	Applies a filter to display only the selected type of items.
Zoom In and Zoom Out	Zoom in and zoom out of the map.
Export	Exports the graph as a .png image or as a PDF.
	Indicates the period over which the measurements are taken.

Network Monitor Report

The App Scope Network Monitor report (**Monitor** > **App Scope** > **Network Monitor**) displays the bandwidth dedicated to different network functions over the specified period of time. Each network function is color-coded as indicated in the legend below the chart. For example, the image below shows application bandwidth for the past 7 days based on session information.

© 2018 Palo Alto Networks, Inc.



The Network Monitor report contains the following buttons and options.

Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Application	Determines the type of item reported: Application, Application Category, Source, or Destination.
Filter	Applies a filter to display only the selected item. **None** displays all entries.
	Determines whether to display session or byte information.
Export	Exports the graph as a .png image or as a PDF.
	Determines whether the information is presented in a stacked column chart or a stacked area chart.
	Indicates the period over which the change measurements are taken.

Traffic Map Report

The App Scope Traffic Map (**Monitor** > **App Scope** > **Traffic Map**) report shows a geographical view of traffic flows according to sessions or flows.

The firewall uses geolocation for creating traffic maps. The firewall is placed at the bottom of the traffic map screen, if you have not specified the geolocation coordinates (**Device** > **Setup** > **Management,** General Settings section) on the firewall.

© 2018 Palo Alto Networks, Inc.



Each traffic type is color-coded as indicated in the legend below the chart. The Traffic Map report contains the following buttons and options.

Buttons	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Incoming threats	Displays incoming threats.
Outgoing threats	Displays outgoing threats.
	Determines whether to display session or byte information.
Zoom In and Zoom Out	Zoom in and zoom out of the map.
Export	Exports the graph as a .png image or as a PDF.
	Indicates the period over which the change measurements are taken.

© 2018 Palo Alto Networks, Inc.

Use the Automated Correlation Engine

The automated correlation engine is an analytics tool that uses the logs on the firewall to detect actionable events on your network. The engine correlates a series of related threat events that, when combined, indicate a likely compromised host on your network or some other higher level conclusion. It pinpoints areas of risk, such as compromised hosts on the network, allows you to assess the risk and take action to prevent exploitation of network resources. The automated correlation engine uses *correlation objects* to analyze the logs for patterns and when a match occurs, it generates a *correlated event*.



The following models support the automated correlation engine:

- *Panorama—M-Series appliances and virtual appliances*
- *PA-7000 Series firewalls*
- *PA-5200 Series firewalls*
- *PA-5000 Series firewalls*
- *PA-3000 Series firewalls*

- Automated Correlation Engine Concepts on page 334
- View the Correlated Objects on page 335
- Interpret Correlated Events on page 335
- Use the Compromised Hosts Widget in the ACC on page 337

Automated Correlation Engine Concepts

The automated correlation engine uses *correlation objects* to analyze the logs for patterns and when a match occurs, it generates a *correlated event*.

- Correlation Object on page 334
- Correlated Events on page 335

Correlation Object

A correlation object is a definition file that specifies patterns to match against, the data sources to use for the lookups, and time period within which to look for these patterns. A pattern is a boolean structure of conditions that queries the following data sources (or logs) on the firewall: application statistics, traffic, traffic summary, threat summary, threat, data filtering, and URL filtering. Each pattern has a severity rating, and a threshold for the number of times the pattern match must occur within a defined time limit to indicate malicious activity. When the match conditions are met, a correlated event is logged.

A correlation object can connect isolated network events and look for patterns that indicate a more significant event. These objects identify suspicious traffic patterns and network anomalies, including suspicious IP activity, known command-and-control activity, known vulnerability exploits, or botnet activity that, when correlated, indicate with a high probability that a host on the network has been compromised. Correlation objects are defined and developed by the Palo Alto Networks Threat Research team, and are delivered with the weekly dynamic updates to the firewall and Panorama. To obtain new correlation objects, the firewall must have a Threat Prevention license. Panorama requires a support license to get the updates.

The patterns defined in a correlation object can be static or dynamic. Correlated objects that include patterns observed in WildFire are dynamic, and can correlate malware patterns detected by WildFire with command-and-control activity initiated by a host that was targeted with the malware on your network or activity seen by a Traps protected endpoint on Panorama. For example, when a host submits a file to the WildFire cloud and the verdict is malicious, the correlation object looks for other hosts or clients on the network that exhibit the same behavior seen in the cloud. If the malware sample had performed a DNS

© 2018 Palo Alto Networks, Inc.

query and browsed to a malware domain, the correlation object will parse the logs for a similar event. When the activity on a host matches the analysis in the cloud, a high severity correlated event is logged.

Correlated Events

A correlated event is logged when the patterns and thresholds defined in a correlation object match the traffic patterns on your network. To Interpret Correlated Events on page 335 and to view a graphical display of the events, see Use the Compromised Hosts Widget in the ACC on page 337.

View the Correlated Objects

You can view the correlation objects that are currently available on the firewall.

STEP 1 | Select **Monitor** > **Automated Correlation Engine** > **Correlation Objects**. All the objects in the list are enabled by default.



STEP 2 | View the details on each correlation object. Each object provides the following information:

- **Name** and **Title**—The name and title indicate the type of activity that the correlation object detects. The name column is hidden from view, by default. To view the definition of the object, unhide the column and click the name link.
- **ID**— A unique number that identifies the correlation object; this column is also hidden by default. The IDs are in the 6000 series.
- **Category**—A classification of the kind of threat or harm posed to the network, user, or host. For now, all the objects identify compromised hosts on the network.
- **State**—Indicates whether the correlation object is enabled (active) or disabled (inactive). All the objects in the list are enabled by default, and are hence active. Because these objects are based on threat intelligence data and are defined by the Palo Alto Networks Threat Research team, keep the objects active in order to track and detect malicious activity on your network.
- **Description**—Specifies the match conditions for which the firewall or Panorama will analyze logs. It describes the sequence of conditions that are matched on to identify acceleration or escalation of malicious activity or suspicious host behavior. For example, the **Compromise Lifecycle** object detects a host involved in a complete attack lifecycle in a three-step escalation that starts with scanning or probing activity, progressing to exploitation, and concluding with network contact to a known malicious domain.

For more information, see Automated Correlation Engine Concepts on page 334 and Use the Automated Correlation Engine on page 334.

Interpret Correlated Events

You can view and analyze the logs generated for each correlated event in the **Monitor** > **Automated Correlation Engine** > **Correlated Events** tab.

© 2018 Palo Alto Networks, Inc.



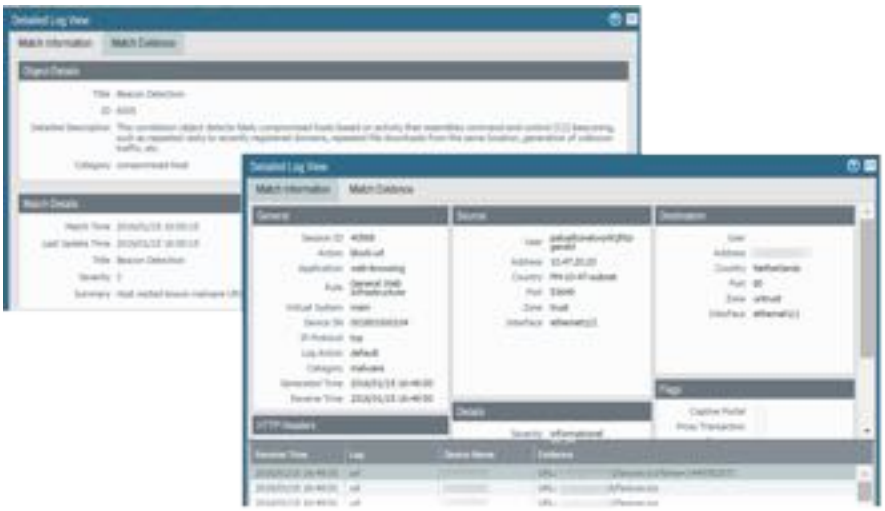
Correlated Events on page 335 includes the following details:

Field	Description
Match Time	The time the correlation object triggered a match.
Update Time	The time when the event was last updated with evidence on the match. As the firewall collects evidence on pattern or sequence of events defined in a correlation object, the time stamp on the correlated event log is updated.
Object Name	The name of the correlation object that triggered the match.
Source Address	The IP address of the user/device on your network from which the traffic originated.
Source User	The user and user group information from the directory server, if User-ID on page 445 is enabled.
Severity *To configure the firewall or Panorama to send alerts using email, SNMP or syslog messages for a desired severity level, see Use External Services for Monitoring on page 377.*	A rating that indicates the urgency and impact of the match. The severity level indicates the extent of damage or escalation pattern, and the frequency of occurrence. Because correlation objects are primarily for detecting threats, the correlated events typically relate to identifying compromised hosts on the network and the severity implies the following: • **Critical**—Confirms that a host has been compromised based on correlated events that indicate an escalation pattern. For example, a critical event is logged when a host that received a file with a malicious verdict by WildFire exhibits the same command-and-control activity that was observed in the WildFire sandbox for that malicious file. • **High**—Indicates that a host is very likely compromised based on a correlation between multiple threat events, such as malware detected anywhere on the network that matches the command-and-control activity generated by a particular host. • **Medium**—Indicates that a host is likely compromised based on the detection of one or multiple suspicious events, such as repeated visits to known malicious URLs, which suggests a scripted command-and-control activity. • **Low**—Indicates that a host is possibly compromised based on the detection of one or multiple suspicious events, such as a visit to a malicious URL or a dynamic DNS domain. • **Informational**—Detects an event that may be useful in aggregate for identifying suspicious activity, but the event is not necessarily significant on its own.

© 2018 Palo Alto Networks, Inc.

Field	Description
Summary	A description that summarizes the evidence gathered on the correlated event.

Click the icon to see the detailed log view, which includes all the evidence on a match:



Tab	Description
Match Information	Object Details: Presents information on the Correlation Object on page 334 that triggered the match.
	Match Details: A summary of the match details that includes the match time, last update time on the match evidence, severity of the event, and an event summary.
Match Evidence	Presents all the evidence that corroborates the correlated event. It lists detailed information on the evidence collected for each session.

Use the Compromised Hosts Widget in the ACC

The compromised hosts widget on **ACC** > **Threat Activity**, aggregates the Correlated Events on page 335 and sorts them by severity. It displays the source IP address/user who triggered the event, the correlation object that was matched and the number of times the object was matched. Use the match count link to jump to the match evidence details.



© 2018 Palo Alto Networks, Inc.

For more details, see Use the Automated Correlation Engine on page 334 and Use the Application Command Center on page 309.

© 2018 Palo Alto Networks, Inc.

Take Packet Captures

All 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 packet captures on the dataplane, you may need to Disable Hardware Offload on page 339 to ensure that the firewall captures all traffic.



Packet capture can be very CPU intensive and can degrade firewall performance. Only use this feature when necessary and make sure you turn it off after you have collected the required packets.

- Types of Packet Captures on page 339
- Disable Hardware Offload on page 339
- Take a Custom Packet Capture on page 340
- Take a Threat Packet Capture on page 344
- Take an Application Packet Capture on page 345
- Take a Packet Capture on the Management Interface on page 348

Types of Packet Captures

There are four different types of packet captures you can enable, depending on what you need to do:

- **Custom Packet Capture**—The firewall captures packets for all traffic or for specific traffic based on 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 then use the packet captures for troubleshooting network#related issues or for gathering application attributes to enable you to write custom application signatures or to request an application signature from Palo Alto Networks. See Take a Custom Packet Capture on page 340.
- **Threat Packet Capture**—The firewall captures packets when it detects a virus, spyware, or vulnerability. You enable this feature in Antivirus, Anti-Spyware, and Vulnerability Protection security profiles. A link to view or export the packet captures will appear 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. See Take a Threat Packet Capture on page 344.
- **Application Packet Capture**—The firewall captures packets based on a specific application and filters that you define. A link to view or export the packet captures will appear in the second column of the Traffic logs for traffic that matches the packet capture rule. See Take an Application Packet Capture on page 345.
- **Management Interface Packet Capture**—The firewall captures packets on the management interface (MGT). The packet captures are useful when troubleshooting services that traverse the interface, such as firewall management authentication to External Authentication Services on page 155, software and content updates, log forwarding, communication with SNMP servers, and authentication requests for GlobalProtect and Captive Portal. See Take a Packet Capture on the Management Interface on page 348.

Disable Hardware Offload

Packet captures for traffic passing through the network data ports on a Palo Alto Networks firewall are performed by the dataplane CPU. To capture traffic that passes through the management interface, you must Take a Packet Capture on the Management Interface, in which case the packet capture is performed on the management plane.

© 2018 Palo Alto Networks, Inc.

When a packet capture is performed on the dataplane, the packet capture filter is used differently by the ingress stage, compared to the firewall, drop, and egress capture stages. The ingress stage uses the packet capture filter to copy individual packets that match the filter to the capture file. Packets that fail packet-parsing checks are dropped before being captured. The firewall, drop, and egress capture stages use the same packet capture filter to mark all new sessions that match the filter. Because each session, as recorded in the session tables, identifies both client-to-server and server-to-client connections, any traffic, in either direction, that matches to the flagged session will be copied to the firewall-stage and transmit-stage capture files. Likewise, any dropped traffic (post receive stage) in either direction that matches to a flagged session will be copied to the drop-stage capture file.

On firewall models that include a network processor, traffic that meets certain pre-determined criteria by Palo Alto Networks may be offloaded for handling by the network processor. Such offloaded traffic will not reach the dataplane CPU and will, therefore, not be captured. To capture offloaded traffic, you must use the CLI to turn off the hardware offload feature.

Common types of traffic that may be offloaded include non-decrypted SSL and SSH traffic (which being encrypted cannot be usefully inspected beyond the initial SSL/SSH session setup), network protocols (such as OSPF, BGP, RIP), and traffic that matches an application-override policy. Some types of traffic will never be offloaded, such as ARP, all non-IP traffic, IPSec, and VPN sessions. Individual SYN, FIN, and RST packets, even for session traffic that has been offloaded, will never be offloaded, and will always be passed through to the dataplane CPU, once recognized as such by the network processor.

Hardware offload is supported on the following firewalls: PA-3050, PA-3060, PA-5000 Series, PA-5200 Series, and PA-7000 Series firewall.

Disabling hardware offload may increase the dataplane CPU usage. If dataplane CPU usage is already high, you may want to schedule a maintenance window before disabling hardware offload.

STEP 1 | Disable hardware offload by running the following CLI command:

```
admin@PA-7050>set session offload no
```

STEP 2 | After the firewall captures the required traffic, enable hardware offload by running the following CLI command:

```
admin@PA-7050>set session offload yes
```

Take a Custom Packet Capture

Custom packet captures allow you to define the traffic that the firewall will capture. To ensure that you capture all traffic, you may need to Disable Hardware Offload on page 339.

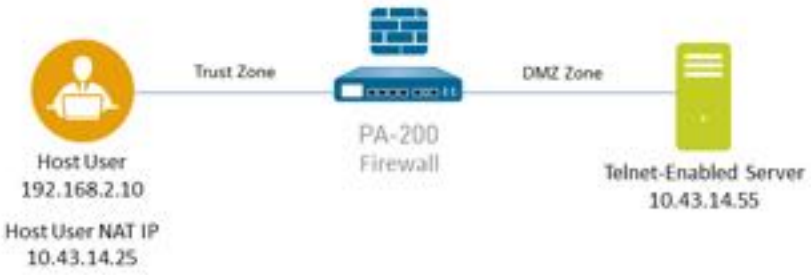
STEP 1 | Before you start a packet capture, identify the attributes of the traffic that you want to capture.

For example, to determine the source IP address, source NAT IP address, and the destination IP address for traffic between two systems, perform a ping from the source system to the to the destination system. After the ping is complete, go to **Monitor** > **Traffic** and locate the traffic log for the two systems. Click the **Detailed Log View** icon located in the first column of the log and note the source address, source NAT IP, and the destination address.

© 2018 Palo Alto Networks, Inc.



The following example shows how to use a packet capture to troubleshoot a Telnet connectivity issue from a user in the Trust zone to a server in the DMZ zone.



STEP 2 | Set packet capture filters, so the firewall only captures traffic you are interested in.

Using filters makes it easier for you to locate the information you need in the packet capture and will reduce the processing power required by the firewall to take the packet capture. To capture all traffic, do not define filters and leave the filter option off.

For example, if you configured NAT on the firewall, you will need to apply two filters. The first one filters on the pre-NAT source IP address to the destination IP address and the second one filters traffic from the destination server to the source NAT IP address.

1. Select **Monitor** > **Packet Capture**.
2. Click **Clear All Settings** at the bottom of the window to clear any existing capture settings.
3. Click **Manage Filters** and click **Add**.
4. Select **Id** 1 and in the **Source** field enter the source IP address you are interested in and in the **Destination** field enter a destination IP address.

   For example, enter the source IP address **`192.168.2.10`** and the destination IP address **`10.43.14.55`**. To further filter the capture, set **Non-IP** to **exclude** non-IP traffic, such as broadcast traffic.
5. **Add** the second filter and select **Id** 2.

   For example, in the **Source** field enter **`10.43.14.55`** and in the **Destination** field enter **`10.43.14.25`**. In the **Non-IP** drop-down menu select **exclude**.



6. Click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Set **Filtering** to **On**.

STEP 4 | Specify the traffic stage(s) that trigger the packet capture and the filename(s) to use to store the captured content. For a definition of each stage, click the **Help** icon on the packet capture page.

For example, to configure all packet capture stages and define a filename for each stage, perform the following procedure:

1. **Add** a **Stage** to the packet capture configuration and define a **File** name for the resulting packet capture.

   For example, select **receive** as the **Stage** and set the **File** name to telnet-test-received.

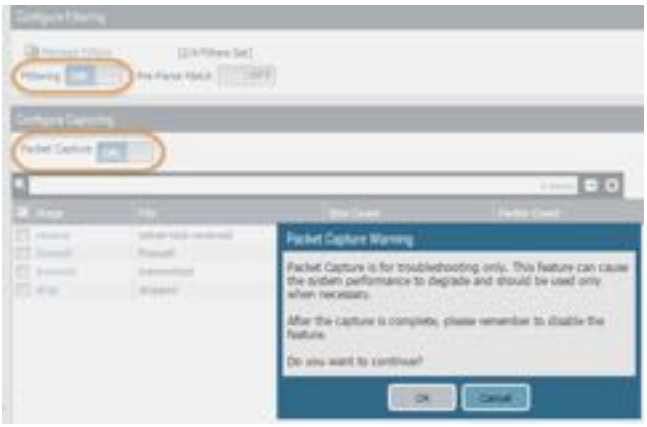


2. Continue to **Add** each **Stage** you want to capture (**receive, firewall**, **transmit**, and **drop**) and set a unique **File** name for each stage.



STEP 5 | Set **Packet Capture** to **ON**.

The firewall or appliance warns you that system performance can be degraded; acknowledge the warning by clicking **OK**. If you define filters, the packet capture should have little impact on performance, but you should always turn **Off** packet capture after the firewall captures the data that you want to analyze.



STEP 6 | Generate traffic that matches the filters that you defined.

For this example, generate traffic from the source system to the Telnet-enabled server by running the following command from the source system (192.168.2.10):

```
telnet 10.43.14.55
```


© 2018 Palo Alto Networks, Inc.

STEP 7 | Turn packet capture **OFF** and then click the refresh icon to see the packet capture files.



Notice that in this case, there were no dropped packets, so the firewall did not create a file for the drop stage.

STEP 8 | Download the packet captures by clicking the filename in the File Name column.



STEP 9 | View the packet capture files using a network packet analyzer.

In this example, the received.pcap packet capture shows a failed Telnet session from the source system at 192.168.2.10 to the Telnet-enabled server at 10.43.14.55. The source system sent the Telnet request to the server, but the server did not respond. In this example, the server may not have Telnet enabled, so check the server.



STEP 10 | Enable the Telnet service on the destination server (10.43.14.55) and turn on packet capture to take a new packet capture.

STEP 11 | Generate traffic that will trigger the packet capture.

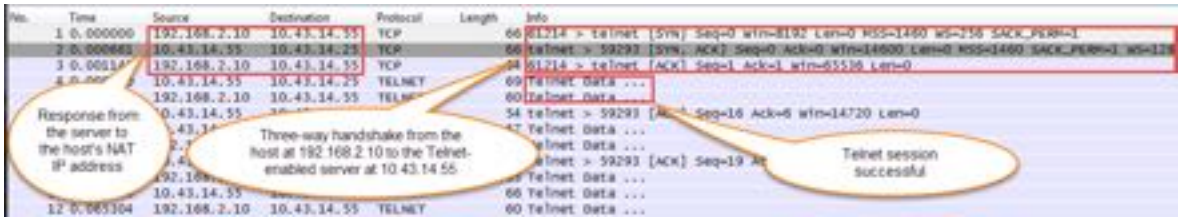
Run the Telnet session again from the source system to the Telnet-enabled server

```
telnet 10.43.14.55
```

STEP 12 | Download and open the received.pcap file and view it using a network packet analyzer.

The following packet capture now shows a successful Telnet session from the host user at 192.168.2.10 to the Telnet-enabled server at 10.43.14.55.

You also see the NAT address 10.43.14.25. When the server responds, it does so to the NAT address. You can see the session is successful as indicated by the three-way handshake between the host and the server and then you see Telnet data.



© 2018 Palo Alto Networks, Inc.

Take a Threat Packet Capture

To configure the firewall to take a packet capture (pcap) when it detects a threat, enable packet capture on Antivirus, Anti-Spyware, and Vulnerability Protection security profiles.

STEP 1 | Enable the packet capture option in the security profile.

Some security profiles allow you to define a single-packet capture, or extended-capture. If you choose extended-capture, define the capture length. This will allow the firewall to capture more packets to provide additional context related to the threat.



If the action for a given threat is set to an action other than allow, the firewall captures only the packet(s) that match the threat signature.

1. Select **Objects** > **Security Profiles** and enable the packet capture option for the supported profiles as follows:
   - **Antivirus**—Select a custom antivirus profile and in the **Antivirus** tab select the **Packet Capture** check box.
   - **Anti-Spyware**—Select a custom Anti-Spyware profile, click the **DNS Signatures** tab and in the **Packet Capture** drop-down, select **single-packet** or **extended-capture**.
   - **Vulnerability Protection**—Select a custom Vulnerability Protection profile and in the **Rules** tab, click **Add** to add a new rule, or select an existing rule. Set **Packet Capture** to **single-packet** or **extended-capture**.

   If the profile has signature exceptions defined, click the Exceptions tab and in the Packet Capture column for a signature, set single-packet or extended-capture.
2. (Optional) If you selected **extended-capture** for any of the profiles, define the extended packet capture length.
   1. Select **Device** > **Setup** > **Content-ID** and edit the Content-ID Settings.
   2. In the **Extended Packet Capture Length (packets)** section, specify the number of packets that the firewall will capture (range is 1-50; default is 5).
   3. Click **OK**.

STEP 2 | Add the security profile (with packet capture enabled) to a Security Policy rule.

1. Select **Policies** > **Security** and select a rule.
2. Select the **Actions** tab.
3. In the Profile Settings section, select a profile that has packet capture enabled.

   For example, click the **Antivirus** drop-down and select a profile that has packet capture enabled.

STEP 3 | View/export the packet capture from the Threat logs.

1. Select **Monitor** > **Logs** > **Threat**.
2. In the log entry that you are interested in, click the green packet capture icon in the second column. View the packet capture directly or **Export** it to your system.

© 2018 Palo Alto Networks, Inc.



Take an Application Packet Capture

The following topics describe two ways that you can configure the firewall to take application packet captures:

- Take a Packet Capture for Unknown Applications on page 345
- Take a Custom Application Packet Capture on page 346

Take a Packet Capture for Unknown Applications

Palo Alto Networks firewalls automatically generate a packet capture for sessions that contain an application that it cannot identify. Typically, the only applications that are classified as unknown traffic—tcp, udp or non-syn-tcp—are commercially available applications that do not yet have App-ID signatures, are internal or custom applications on your network, or potential threats. You can use these packet captures to gather more context related to the unknown application or use the information to analyze the traffic for potential threats. You can also Manage Custom or Unknown Applications by controlling them through security policy or by writing a custom application signature and creating a security rule based on the custom signature. If the application is a commercial application, you can submit the packet capture to Palo Alto Networks to have an App-ID signature created.

STEP 1 | Verify that unknown application packet capture is enabled. This option is on by default.

1. To view the unknown application capture setting, run the following CLI command:

```
admin@PA-200>show running application setting | match "Unknown capture"
```

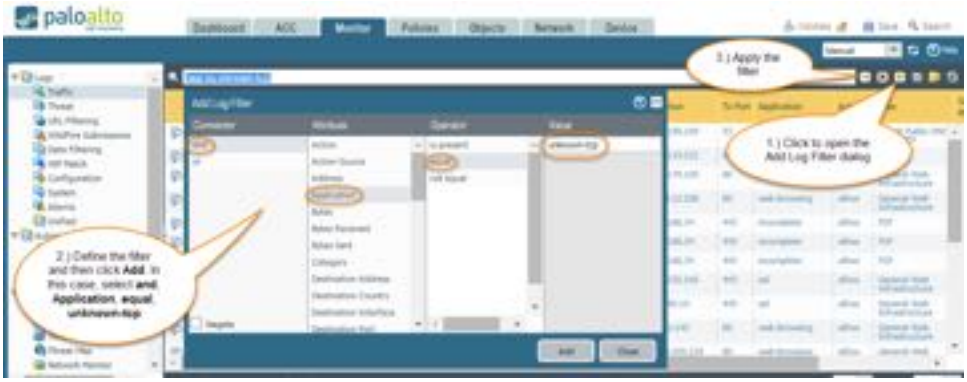
2. If the unknown capture setting option is off, enable it:

```
admin@PA-200>set application dump-unknown yes
```

STEP 2 | Locate unknown application by filtering the traffic logs.

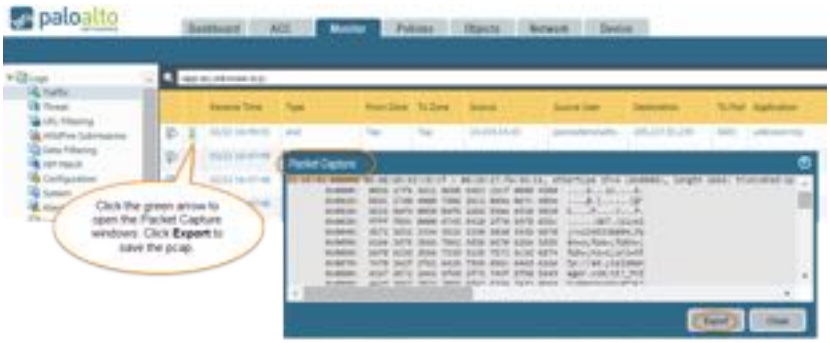
1. Select **Monitor** > **Logs** > **Traffic**.
2. Click **Add Filter** and select the filters as shown in the following example.

© 2018 Palo Alto Networks, Inc.



3. Click **Add** and **Apply Filter**.

STEP 3 | Click the packet capture icon to view the packet capture or **Export** it to your local system.



Take a Custom Application Packet Capture

You can configure a Palo Alto Networks firewall to take a packet capture based on an application name and filters that you define. You can then use the packet capture to troubleshoot issues with controlling an application. When configuring an application packet capture, you must use the application name defined in the App-ID database. You can view a list of all App-ID applications using Applipedia or from the web interface on the firewall in **Objects** > **Applications**.

STEP 1 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall.

STEP 2 | Turn on the application packet capture and define filters.

```
admin@PA-200>set application dump on application <application-name>
 rule <rule-name>
```

For example, to capture packets for the facebook-base application that matches the security rule named rule1, run the following CLI command:

```
admin@PA-200>set application dump on application facebook-base rule rule1
```

You can also apply other filters, such as source IP address and destination IP address.

STEP 3 | View the output of the packet capture settings to ensure that the correct filters are applied. The output appears after enabling the packet capture.

© 2018 Palo Alto Networks, Inc.

In the following output, you see that application filtering is now on based on the facebook-base application for traffic that matches rule1.

```
Application setting:
 Application cache : yes
 Supernode : yes
 Heuristics : yes
 Cache Threshold : 16
 Bypass when exceeds queue limit: no
 Traceroute appid : yes
 Traceroute TTL threshold : 30
 Use cache for appid : no
 Unknown capture : on
 Max. unknown sessions : 5000
 Current unknown sessions : 0
 Application capture : on
 Max. application sessions : 5000
 Current application sessions : 0
 Application filter setting:
 Rule : rule1
 From : any
 To : any
 Source : any
 Destination : any
 Protocol : any
 Source Port : any
 Dest. Port : any
 Application : facebook-base
 Current APPID Signature
 Signature Usage : 21 MB (Max. 32 MB)
 TCP 1 C2S : 15503 states
 TCP 1 S2C : 5070 states
 TCP 2 C2S : 2426 states
 TCP 2 S2C : 702 states
 UDP 1 C2S : 11379 states
 UDP 1 S2C : 2967 states
 UDP 2 C2S : 755 states
 UDP 2 S2C : 224 states
```

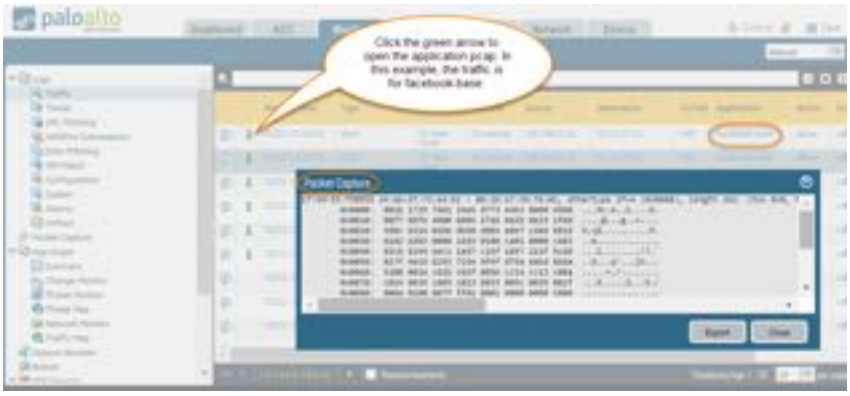
STEP 4 | Access Facebook.com from a web browser to generate Facebook traffic and then turn off application packet capture by running the following CLI command:

```
admin@PA-200>set application dump off
```

STEP 5 | View/export the packet capture.

1. Log in to the web interface on the firewall and select **Monitor** > **Logs** > **Traffic**.
2. In the log entry that you are interested in, click the green packet capture icon in the second column.
3. View the packet capture directly or **Export** it to your computer. The following screen capture shows the facebook-base packet capture.

© 2018 Palo Alto Networks, Inc.



Take a Packet Capture on the Management Interface

The **tcpdump** CLI command enables you to capture packets that traverse the management interface (MGT) on a Palo Alto Networks firewall.

*Each platform has a default number of bytes that **tcpdump** captures. The PA-200 and PA-500 firewalls capture 68 bytes of data from each packet and anything over that is truncated. The PA-3000, PA-5000 Series, the PA-7000 Series firewalls, and VM-Series firewalls capture 96 bytes of data from each packet. To define the number of packets that **tcpdump** will capture, use the **snaplen** (snap length) option (range 0-65535). Setting the **snaplen** to 0 will cause the firewall to use the maximum length required to capture whole packets.*

STEP 1 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall.

STEP 2 | To start a packet capture on the MGT interface, run the following command:

```
admin@PA-200>tcpdump filter "<filter-option> <IP-address>" snaplen length
```

For example, to capture the traffic that is generated when and administrator authenticates to the firewall using RADIUS, filter on the destination IP address of the RADIUS server (10.5.104.99 in this example):

```
admin@PA-200>tcpdump filter "dst 10.5.104.99" snaplen 0
```

You can also filter on src (source IP address), host, net, and you can exclude content. For example, to filter on a subnet and exclude all SCP, SFTP, and SSH traffic (which uses port 22), run the following command:

```
admin@PA-200>tcpdump filter "net 10.5.104.0/24 and not port 22" snaplen 0
```

*Each time **tcpdump** takes a packet capture, it stores the content in a file named mgmt.pcap. This file is overwritten each time you run **tcpdump**.*

STEP 3 | After the traffic you are interested in has traversed the MGT interface, press Ctrl + C to stop the capture.

STEP 4 | View the packet capture by running the following command:

© 2018 Palo Alto Networks, Inc.

```
admin@PA-200> view-pcap mgmt-pcap mgmt.pcap
```

The following output shows the packet capture from the MGT port (10.5.104.98) to the RADIUS server (10.5.104.99):

```
09:55:29.139394 IP 10.5.104.98.43063 > 10.5.104.99.radius: RADIUS, Access
 Request (1), id: 0x00 length: 89
 09:55:29.144354 arp reply 10.5.104.98 is-at 00:25:90:23:94:98 (oui
 Unknown)
 09:55:29.379290 IP 10.5.104.98.43063 > 10.5.104.99.radius: RADIUS, Access
 Request (1), id: 0x00 length: 70
 09:55:34.379262 arp who-has 10.5.104.99 tell 10.5.104.98
```

STEP 5 | (Optional) Export the packet capture from the firewall using SCP (or TFTP). For example, to export the packet capture using SCP, run the following command:

```
admin@PA-200>scp export mgmt-pcap from mgmt.pcap to <username@host:path>
```

For example, to export the pcap to an SCP enabled server at 10.5.5.20 to a temp folder named temp-SCP, run the following CLI command:

```
admin@PA-200>scp export mgmt-pcap from mgmt.pcap to admin@10.5.5.20:c:/temp-
SCP
```

Enter the login name and password for the account on the SCP server to enable the firewall to copy the packet capture to the c:\temp-SCP folder on the SCP-enabled.

STEP 6 | You can now view the packet capture files using a network packet analyzer, such as Wireshark.

© 2018 Palo Alto Networks, Inc.

Monitor Applications and Threats

All Palo Alto Networks next-generation firewalls come equipped with the App-ID on page 509 technology, which identifies the applications traversing your network, irrespective of protocol, encryption, or evasive tactic. You can then Use the Application Command Center on page 309 to monitor the applications. The ACC graphically summarizes the data from a variety of log databases to highlight the applications traversing your network, who is using them, and their potential security impact. ACC is dynamically updated, using the continuous traffic classification that App-ID performs; if an application changes ports or behavior, App-ID continues to see the traffic, displaying the results in ACC. Additional visibility into URL categories, threats, and data provides a complete and well-rounded picture of network activity. With ACC, you can very quickly learn more about the traffic traversing the network and then translate that information into a more informed security policy

You can also Use the Dashboard on page 307 to monitor the network.



Content Delivery Network Infrastructure for Dynamic Updates on page 586 to check whether logged events on the firewall pose a security risk. The AutoFocus intelligence summary shows the prevalence of properties, activities, or behaviors associated with logs in your network and on a global scale, as well as the WildFire verdict and AutoFocus tags linked to them. With an active AutoFocus subscription, you can use this information to create customized AutoFocus Alerts that track specific threats on your network.

© 2018 Palo Alto Networks, Inc.

View and Manage Logs

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.

- Log Types and Severity Levels on page 351
- View Logs on page 356
- Filter Logs on page 357
- Export Logs on page 357
- Configure Log Storage Quotas and Expiration Periods on page 358
- Schedule Log Exports to an SCP or FTP Server on page 358

Log Types and Severity Levels

You can see the following log types in the **Monitor** > **Logs** pages.

- Traffic Logs on page 351
- Threat Logs on page 352
- URL Filtering Logs on page 352
- WildFire Submissions Logs on page 353
- Data Filtering Logs on page 353
- Correlation Logs on page 353
- Tunnel Inspection Logs on page 354
- Config Logs on page 354
- System Logs on page 354
- HIP Match Logs on page 355
- User-ID Logs on page 355
- Alarms Logs on page 355
- Authentication Logs on page 355
- Unified Logs on page 355

Traffic Logs

Traffic logs display an entry for the start and end of each session. Each entry includes the following information: date and time; source and destination zones, addresses and ports; application name; security rule applied to the traffic flow; rule action (allow, deny, or drop); ingress and egress interface; number of bytes; and session end reason.

The Type column indicates whether the entry is for the start or end of the session. The Action column indicates whether the firewall allowed, denied, or dropped the session. A drop indicates the security rule that blocked the traffic specified any application, while a deny indicates the rule identified a specific application. If the firewall drops traffic before identifying the application, such as when a rule drops all traffic for a specific service, the Application column displays not-applicable.

Click beside an entry to view additional details about the session, such as whether an ICMP entry aggregates multiple sessions between the same source and destination (in which case the Count column value is greater than one).

© 2018 Palo Alto Networks, Inc.

Threat Logs

Threat logs display entries when traffic matches one of the Security Profiles attached to a security rule on the firewall. Each entry includes the following information: date and time; type of threat (such as virus or spyware); threat description or URL (Name column); source and destination zones, addresses, and ports; application name; alarm action (such as allow or block); and severity level.

To see more details on individual Threat log entries:

- Click beside a threat entry to view details such as whether the entry aggregates multiple threats of the same type between the same source and destination (in which case the Count column value is greater than one).
- If you configured the firewall to Take Packet Captures, click beside an entry to access the captured packets.

The following table summarizes the Threat severity levels:

Severity	Description
Critical	Serious threats, such as those that affect default installations of widely deployed software, result in root compromise of servers, and the exploit code is widely available to attackers. The attacker usually does not need any special authentication credentials or knowledge about the individual victims and the target does not need to be manipulated into performing any special functions.
High	Threats that have the ability to become critical but have mitigating factors; for example, they may be difficult to exploit, do not result in elevated privileges, or do not have a large victim pool. WildFire Submissions log entries with a malicious verdict and an action set to allow are logged as High.
Medium	Minor threats in which impact is minimized, such as DoS attacks that do not compromise the target or exploits that require an attacker to reside on the same LAN as the victim, affect only non-standard configurations or obscure applications, or provide very limited access. WildFire Submissions log entries with a grayware verdict are logged as Medium.
Low	Warning-level threats that have very little impact on an organization's infrastructure. They usually require local or physical system access and may often result in victim privacy or DoS issues and information leakage. Data Filtering profile matches are logged as Low. WildFire Submissions log entries with a grayware verdict and an action set to allow are logged as Low.
Informational	Suspicious events that do not pose an immediate threat, but that are reported to call attention to deeper problems that could possibly exist. URL Filtering log entries are logged as Informational. WildFire Submissions log entries with a benign verdict and an action set to allow are logged as Informational. Additionally, log entries with any verdict and an action set to block are also logged as Informational.

URL Filtering Logs

URL Filtering on page 641 logs display entries for traffic that matches URL Filtering Profile on page 647 attached to security rules. For example, the firewall generates a log if a rule blocks access to specific web sites and web site categories or if you configured a rule to generate an alert when a user accesses a web site.

© 2018 Palo Alto Networks, Inc.

WildFire Submissions Logs

The firewall forwards samples (files and emails links) to the WildFire cloud for analysis based on WildFire Analysis profiles settings (**Objects** > **Security Profiles** > **WildFire Analysis**). The firewall generates WildFire Submissions log entries for each sample it forwards after WildFire completes 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.

The following table summarizes the WildFire verdicts:

Severity	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 an 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 are can pose a security threat. Malware can include viruses, 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.

Data Filtering Logs

Data Filtering logs display entries for the security rules that help prevent sensitive information such as credit card numbers from leaving the area that the firewall protects. See Set Up Data Filtering on page 538 for information on defining Data Filtering profiles.

This log type also shows information for File Blocking Profiles. For example, if a rule blocks .exe files, the log shows the blocked files.

Correlation Logs

The firewall logs a correlated event when the patterns and thresholds defined in a Correlation Object on page 334 match the traffic patterns on your network. To Interpret Correlated Events on page 335 and view a graphical display of the events, see Use the Compromised Hosts Widget in the ACC on page 337.

The following table summarizes the Correlation log severity levels:

Severity	Description
Critical	Confirms that a host has been compromised based on correlated events that indicate an escalation pattern. For example, a critical event is logged when a host that received a file with a malicious verdict by WildFire, exhibits the same command-and control activity that was observed in the WildFire sandbox for that malicious file.

© 2018 Palo Alto Networks, Inc.

Severity	Description
High	Indicates that a host is very likely compromised based on a correlation between multiple threat events, such as malware detected anywhere on the network that matches the command and control activity being generated from a particular host.
Medium	Indicates that a host is likely compromised based on the detection of one or multiple suspicious events, such as repeated visits to known malicious URLs that suggests a scripted command-and-control activity.
Low	Indicates that a host is possibly compromised based on the detection of one or multiple suspicious events, such as a visit to a malicious URL or a dynamic DNS domain.
Informational	Detects an event that may be useful in aggregate for identifying suspicious activity; each event is not necessarily significant on its own.

Tunnel Inspection Logs

Tunnel inspection logs are like traffic logs for tunnel sessions; they display entries of non-encrypted tunnel sessions. To prevent double counting, the firewall saves only the inner flows in traffic logs, and sends tunnel sessions to the tunnel inspection logs. The tunnel inspection log entries include Receive Time (date and time the log was received), the tunnel ID, monitor tag, session ID, the Security rule applied to the tunnel session, number of bytes in the session, parent session ID (session ID for the tunnel session), source address, source user and source zone, destination address, destination user, and destination zone.

Click the Detailed Log view to see details for an entry, such as the tunnel protocol used, and the flag indicating whether the tunnel content was inspected or not. Only a session that has a parent session will have the Tunnel Inspected flag set, which means the session is in a tunnel-in-tunnel (two levels of encapsulation). The first outer header of a tunnel will not have the Tunnel Inspected flag set.

Config Logs

Config logs display entries for changes to the firewall configuration. Each entry includes the date and time, the administrator username, the IP address from where the administrator made the change, the type of client (Web, CLI, or Panorama), the type of command executed, the command status (succeeded or failed), the configuration path, and the values before and after the change.

System Logs

System logs display entries for each system event on the firewall. Each entry includes the date and time, event severity, and event description. The following table summarizes the System log severity levels. For a partial list of System log messages and their corresponding severity levels, refer to System Log Events.

Severity	Description
Critical	Hardware failures, including high availability (HA) failover and link failures.
High	Serious issues, including dropped connections with external devices, such as LDAP and RADIUS servers.
Medium	Mid-level notifications, such as antivirus package upgrades.
Low	Minor severity notifications, such as user password changes.

© 2018 Palo Alto Networks, Inc.

Severity	Description
Informational	Log in/log off, administrator name or password change, any configuration change, and all other events not covered by the other severity levels.

HIP Match Logs

The GlobalProtect Host Information Profile (HIP) matching enables you to collect information about the security status of the end devices accessing your network (such as whether they have disk encryption enabled). The firewall can allow or deny access to a specific host based on adherence to the HIP-based security rules you define. HIP Match logs display traffic flows that match a HIP Object or HIP Profile that you configured for the rules.

User-ID Logs

User-ID on page 445 logs display information about IP address-to-username mappings and Authentication Timestamps on page 211, such as the sources of the mapping information and the times when users authenticated. You can use this information to help troubleshoot User-ID and authentication issues. For example, if the firewall is applying the wrong policy rule for a user, you can view the logs to verify whether that user is mapped to the correct IP address and whether the group associations are correct.

Alarms Logs

An alarm is a firewall-generated message indicating that the number of events of a particular type (for example, encryption and decryption failures) has exceeded the threshold configured for that event type. To enable alarms and configure alarm thresholds, select **Device** > **Log Settings** and edit the Alarm Settings.

When generating an alarm, the firewall creates an Alarm log and opens the System Alarms dialog to display the alarm. After you **Close** the dialog, you can reopen it anytime by clicking **Alarms** ( ) at the bottom of the web interface. To prevent the firewall from automatically opening the dialog for a particular alarm, select the alarm in the Unacknowledged Alarms list and **Acknowledge** the alarm.

Authentication Logs

Authentication logs display information about authentication events that occur when end users try to access network resources for which access is controlled by Authentication Policy on page 211 rules. You can use this information to help troubleshoot access issues and to adjust your Authentication policy as needed. In conjunction with correlation objects, you can also use Authentication logs to identify suspicious activity on your network, such as brute force attacks.

Optionally, you can configure Authentication rules to log timeout events. These timeouts relate to the period when a user need authenticate for a resource only once but can access it repeatedly. Seeing information about the timeouts helps you decide if and how to adjust them (for details, see Authentication Timestamps on page 211).

System logs record authentication events relating to GlobalProtect and to administrator access to the web interface.

Unified Logs

Unified logs are entries from the Traffic, Threat, URL Filtering, WildFire Submissions, and Data Filtering logs displayed in a single view. Unified log view enables you to investigate and filter the latest entries from different log types in one place, instead of searching through each log type separately. Click Effective Queries ( ) in the filter area to select which log types will display entries in Unified log view.

© 2018 Palo Alto Networks, Inc.

The Unified log view displays only entries from logs that you have permission to see. For example, an administrator who does not have permission to view WildFire Submissions logs will not see WildFire Submissions log entries when viewing Unified logs. Administrative Role Types on page 73 define these permissions.

When you Set Up Remote Search in AutoFocus to perform a targeted search on the firewall, the search results are displayed in Unified log view.

View Logs

You can view the different log types on the firewall in a tabular format. The firewall locally stores all log files and automatically generates Configuration and System logs by default. To learn more about the security rules that trigger the creation of entries for the other types of logs, see Log Types and Severity Levels on page 351.

To configure the firewall to forward logs as syslog messages, email notifications, or Simple Network Management Protocol (SNMP) traps, Use External Services for Monitoring on page 377.

STEP 1 | Select a log type to view.

1. Select **Monitor** > **Logs**.
2. Select a log type from the list.

The firewall displays only the logs you have permission to see. For example, if your administrative account does not have permission to view WildFire Submissions logs, the firewall does not display that log type when you access the logs pages. Administrative Role Types on page 73 define the permissions.

STEP 2 | (Optional) Customize the log column display.

1. Click the arrow to the right of any column header, and select **Columns**.
2. Select columns to display from the list. The log updates automatically to match your selections.

STEP 3 | View additional details about log entries.

- Click the spyglass ( ) for a specific log entry. The Detailed Log View has more information about the source and destination of the session, as well as a list of sessions related to the log entry.
- (Threat log only) Click next to an entry to access local packet captures of the threat. To enable local packet captures, see Take Packet Captures on page 339.
- (Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering, and Unified logs only) View AutoFocus threat data for a log entry.
  1. Enable AutoFocus Threat Intelligence on page 48.

     Enable AutoFocus in Panorama to view AutoFocus threat data for all Panorama log entries, including those from firewalls that are not connected to AutoFocus and/or are running PAN-OS 7.0 and earlier release versions (Panorama > Setup > Management > AutoFocus).
  2. Hover over an IP address, URL, user agent, threat name (subtype: virus and wildfire-virus only), filename, or SHA-256 hash.
  3. Click the drop-down ( ) and select **AutoFocus**.
  4. Content Delivery Network Infrastructure for Dynamic Updates on page 586.

Next Steps...

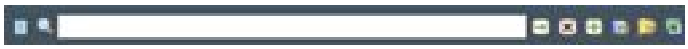
- Filter Logs on page 357.

© 2018 Palo Alto Networks, Inc.

- Export Logs on page 357.
- Configure Log Storage Quotas and Expiration Periods on page 358.

Filter Logs

Each log has a filter area that allows you to set a criteria for which log entries to display. The ability to filter logs is useful for focusing on events on your firewall that possess particular properties or attributes. Filter logs by artifacts that are associated with individual log entries.



STEP 1 | (Unified logs only) Select the log types to include in the Unified log display.

1. Click Effective Queries ( ).
2. Select one or more log types from the list (**traffic**, **threat**, **url**, **data**, **wildfire**, **tunnel**, or **auth**).
3. Click **OK**. The Unified log updates to show only entries from the log types you have selected.

STEP 2 | Add a filter to the filter field.

If the value of the artifact matches the operator (such as has or in), enclose the value in quotation marks to avoid a syntax error. For example, if you filter by destination country and use IN as a value to specify INDIA, enter the filter as ***(dstloc eq "IN")***.

- Click one or more artifacts (such as the application type associated with traffic and the IP address of an attacker) in a log entry. For example, click the Source **10.0.0.25** and Application **web-browsing** of a log entry to display only entries that contain both artifacts in the log (AND search).
- To specify artifacts to add to the filter field, click Add Filter ( ).
- To add a previously saved filter, click Load Filter ( ).

STEP 3 | Apply the filter to the log.

Click Apply Filter ( ). The log will refresh to display only log entries that match the current filter.

STEP 4 | (Optional) Save frequently used filters.

1. Click Save Filter ( ).
2. Enter a **Name** for the filter.
3. Click **OK**. You can view your saved filters by clicking Load Filter ( ).

Next Steps...

- View Logs.
- Export Logs.

Export Logs

You can export the contents of a log type to a comma-separated value (CSV) formatted report. By default, the report contains up to 2,000 rows of log entries.

STEP 1 | Set the number of rows to display in the report.

1. Select **Device** > **Setup** > **Management**, then edit the Logging and Reporting Settings.
2. Click the **Log Export and Reporting** tab.
3. Edit the number of **Max Rows in CSV Export** (up to 1048576 rows).
4. Click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 2 | Download the log.

1. Click Export to CSV ( ). A progress bar showing the status of the download appears.
2. When the download is complete, click **Download file** to save a copy of the log to your local folder. For descriptions of the column headers in a downloaded log, refer to Syslog Field Descriptions.

Next Step...

Schedule Log Exports to an SCP or FTP Server.

Configure Log Storage Quotas and Expiration Periods

The firewall automatically deletes logs that exceed the expiration period. When the firewall reaches the storage quota for a log type, it automatically deletes older logs of that type to create space even if you don't set an expiration period.

If you want to manually delete logs, select Device > Log Settings and, in the Manage Logs section, click the links to clear logs by type.

STEP 1 | Select **Device** > **Setup** > **Management** and edit the Logging and Reporting Settings.

STEP 2 | Select **Log Storage** and enter a **Quota (%)** for each log type. When you change a percentage value, the dialog refreshes to display the corresponding absolute value (Quota GB/MB column).

STEP 3 | Enter the **Max Days** (expiration period) for each log type (range is 1-2,000). The fields are blank by default, which means the logs never expire.

The firewall synchronizes expiration periods across high availability (HA) pairs. Because only the active HA peer generates logs, the passive peer has no logs to delete unless failover occurs and it starts generating logs.

STEP 4 | Click **OK** and **Commit**.

Schedule Log Exports to an SCP or FTP Server

You can schedule exports of Traffic, Threat, URL Filtering, Data Filtering, HIP Match, and WildFire Submission logs to a Secure Copy (SCP) server or File Transfer Protocol (FTP) server. Perform this task for each log type you want to export.

You can use Secure Copy (SCP) commands from the CLI to export the entire log database to an SCP server and import it to another firewall. Because the log database is too large for an export or import to be practical on the following platforms, they do not support these options: PA-7000 Series firewalls (all PAN-OS releases), Panorama virtual appliance running Panorama 6.0 or later releases, and Panorama M-Series appliances (all Panorama releases).

STEP 1 | Select **Device** > **Scheduled Log Export** and click **Add**.

STEP 2 | Enter a **Name** for the scheduled log export and **Enable** it.

STEP 3 | Select the **Log Type** to export.

© 2018 Palo Alto Networks, Inc.

STEP 4 | Select the daily **Scheduled Export Start Time**. The options are in 15-minute increments for a 24-hour clock (00:00 - 23:59).

STEP 5 | Select the **Protocol** to export the logs: **SCP** (secure) or **FTP**.

STEP 6 | Enter the **Hostname** or IP address of the server.

STEP 7 | Enter the **Port** number. By default, FTP uses port 21 and SCP uses port 22.

STEP 8 | Enter the **Path** or directory in which to save the exported logs.

STEP 9 | Enter the **Username** and, if necessary, the **Password** (and **Confirm Password**) to access the server.

STEP 10 | **(**FTP only**)** Select **Enable FTP Passive Mode** if you want to use FTP passive mode, in which the firewall initiates a data connection with the FTP server. By default, the firewall uses FTP active mode, in which the FTP server initiates a data connection with the firewall. Choose the mode based on what your FTP server supports and on your network requirements.

STEP 11 | **(**SCP only**)** Click **Test SCP server connection**. Before establishing a connection, the firewall must accept the host key for the SCP server.



If you use a Panorama template to configure the log export schedule, you must perform this step after committing the template configuration to the firewalls. After the template commit, log in to each firewall, open the log export schedule, and click Test SCP server connection.

STEP 12 | Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Monitor Block List

There are two ways you can cause the firewall to place an IP address on the block list:

- Configure a Vulnerability Protection profile with a rule to Block IP connections and apply the profile to a Security policy, which you apply to a zone.
- Configure a DoS Protection policy rule with the Protect action and a Classified DoS Protection profile, which specifies a maximum rate of connections per second allowed. When incoming packets match the DoS Protection policy and exceed the Max Rate, and if you specified a Block Duration and a Classified policy rule to include source IP address, the firewall puts the offending source IP address on the block list.

In the cases described above, the firewall automatically blocks that traffic in hardware before those packets use CPU or packet buffer resources. If attack traffic exceeds the blocking capacity of the hardware, the firewall uses IP blocking mechanisms in software to block the traffic.

The firewall automatically creates a hardware block list entry based on your Vulnerability Protection profile or DoS Protection policy rule; the source address from the rule is the source IP address in the hardware block list.

Entries on the block list indicate in the Type column whether they were blocked by hardware (hw) or software (sw). The bottom of the screen displays:

- Count of **Total Blocked IPs** out of the number of blocked IP addresses the firewall supports.
- Percentage of the block list that the firewall has used.

To view details about an address on the block list, hover over a Source IP address and click the down arrow link. Click the Who Is link, which displays the Network Solutions Who Is feature, providing information about the address.

For information on configuring a Vulnerability Protection profile, see Customize the Action and Trigger Conditions for a Brute Force Signature. For more information on block list and DoS Protection profiles, see DoS Protection Against Flooding of New Sessions.

© 2018 Palo Alto Networks, Inc.

View and Manage Reports

The reporting capabilities on the firewall allow you to keep a pulse on your network, validate your policies, and focus your efforts on maintaining network security for keeping your users safe and productive.

- Report Types on page 361
- View Reports on page 361
- Configure the Expiration Period and Run Time for Reports on page 362
- Disable Predefined Reports on page 362
- Custom Reports on page 363
- Generate Custom Reports on page 365
- Generate Botnet Reports on page 367
- Generate the SaaS Application Usage Report on page 368
- Manage PDF Summary Reports on page 371
- Generate User/Group Activity Reports on page 373
- Manage Report Groups on page 374
- Schedule Reports for Email Delivery on page 375

Report Types

The firewall includes predefined reports that you can use as-is, or you can build custom reports that meet your needs for specific data and actionable tasks, or you can combine predefined and custom reports to compile information you need. The firewall provides the following types of reports:

- **Predefined Reports**—Allow you to view a quick summary of the traffic on your network. A suite of predefined reports are available in four categories—Applications, Traffic, Threat, and URL Filtering. See View Reports on page 361.
- **User or Group Activity Reports**—Allow you to schedule or create an on-demand report on the application use and URL activity for a specific user or for a user group. The report includes the URL categories and an estimated browse time calculation for individual users. See Generate User/Group Activity Reports on page 373.
- **Custom Reports**—Create and schedule custom reports that show exactly the information you want to see by filtering on conditions and columns to include. You can also include query builders for more specific drill down on report data. See Generate Custom Reports on page 365.
- **PDF Summary Reports**—Aggregate up to 18 predefined or custom reports/graphs from Threat, Application, Trend, Traffic, and URL Filtering categories into one PDF document. See Manage PDF Summary Reports on page 371.
- **Botnet Reports**—Allow you to use behavior-based mechanisms to identify potential botnet-infected hosts in the network. See Generate Botnet Reports on page 367.
- **Report Groups**—Combine custom and predefined reports into report groups and compile a single PDF that is emailed to one or more recipients. See Manage Report Groups on page 374.

Reports can be generated on demand, on a recurring schedule, and can be scheduled for email delivery.

View Reports

The firewall provides an assortment of over 40 predefined reports that it generates every day. You can view these reports directly on the firewall. You can also view custom reports and summary reports.

About 200 MB of storage is allocated for saving reports on the firewall. You can't configure this limit but you can Configure the Expiration Period and Run Time for Reports on page 362 to allow the firewall to delete reports that exceed the period. Keep in mind that when the firewall reaches its storage limit, it

© 2018 Palo Alto Networks, Inc.

automatically deletes older reports to create space even if you don't set an expiration period. Another way to conserve system resources on the firewall is to Disable Predefined Reports on page 362. For long-term retention of reports, you can export the reports (as described below) or Schedule Reports for Email Delivery on page 375.



Unlike other reports, you can't save User/Group Activity reports on the firewall. You must Generate User/Group Activity Reports on page 373 on demand or schedule them for email delivery.

STEP 1 | Select **Monitor** > **Reports**.

The reports are grouped into sections (types) on the right-hand side of the page: **Custom Reports**, **Application Reports**, **Traffic Reports**, **Threat Reports**, **URL Filtering Reports**, and **PDF Summary Reports**.

STEP 2 | Select a report to view. The reports page then displays the report for the previous day.

To view reports for other days, select a date in the calendar at the bottom right of the page and select a report. If you select a report in another section, the date selection resets to the current date.

STEP 3 | To view a report offline, you can export the report to PDF, CSV or to XML formats. Click **Export to PDF**, **Export to CSV**, or **Export to XML** at the bottom of the page, then print or save the file.

Configure the Expiration Period and Run Time for Reports

The expiration period and run time are global settings that apply to all Report Types on page 361. After running new reports, the firewall automatically deletes reports that exceed the expiration period.

STEP 1 | Select **Device** > **Setup** > **Management**, edit the Logging and Reporting Settings, and select the **Log Export and Reporting** tab.

STEP 2 | Set the **Report Runtime** to an hour in the 24-hour clock schedule (default is 02:00; range is 00:00 [midnight] to 23:00).

STEP 3 | Enter the **Report Expiration Period** in days (default is no expiration; range is 1 is 2,000).



You can't change the storage that the firewall allocates for saving reports: it is predefined at about 200 MB. When the firewall reaches the storage maximum, it automatically deletes older reports to create space even if you don't set a Report Expiration Period.

STEP 4 | Click **OK** and **Commit**.

Disable Predefined Reports

The firewall includes about 40 predefined reports that it automatically generates daily. If you do not use some or all of these, you can disable selected reports to conserve system resources on the firewall.

Make sure that no report group or PDF summary report includes the predefined reports you will disable. Otherwise, the firewall will render the PDF summary report or report group without any data.

STEP 1 | Select **Device** > **Setup** > **Management** and edit the Logging and Reporting Settings.

STEP 2 | Select the **Pre-Defined Reports** tab and clear the check box for each report you want to disable. To disable all predefined reports, click **Deselect All**.

© 2018 Palo Alto Networks, Inc.

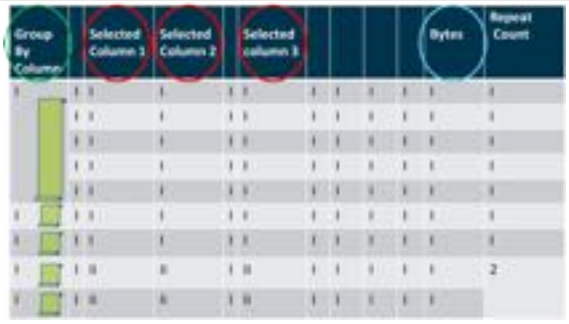
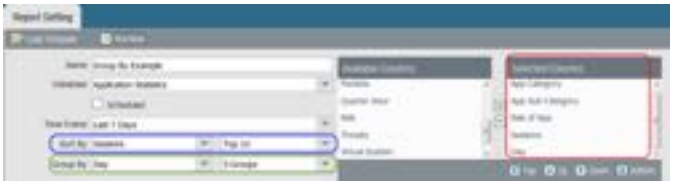
STEP 3 | Click **OK** and **Commit**.

Custom Reports

In order to create purposeful custom reports, you must consider the attributes or key pieces of information that you want to retrieve and analyze. This consideration guides you in making the following selections in a custom report:

Selection	Description
Database	You can base the report on one of the following database types: • Summary databases—These databases are available for Application Statistics, Traffic, Threat, URL Filtering, and Tunnel Inspection logs. The firewall aggregates the detailed logs at 15-minute intervals. To enable faster response time when generating reports, the firewall condenses the data: duplicate sessions are grouped and incremented with a repeat counter, and some attributes (columns) are excluded from the summary. • Detailed logs—These databases itemize the logs and list all the attributes (columns) for each log entry.  *Reports based on detailed logs take much longer to run and are not recommended unless absolutely necessary.*
Attributes	The columns that you want to use as the match criteria. The attributes are the columns that are available for selection in a report. From the list of **Available Columns**, you can add the selection criteria for matching data and for aggregating the details (the **Selected Columns**).
Sort By/ Group By	The **Sort By** and the **Group By** criteria allow you to organize/segment the data in the report; the sorting and grouping attributes available vary based on the selected data source. The Sort By option specifies the attribute that is used for aggregation. If you do not select an attribute to sort by, the report will return the first N number of results without any aggregation. The Group By option allows you to select an attribute and use it as an anchor for grouping data; all the data in the report is then presented in a set of top 5, 10, 25 or 50 groups. For example, when you select Hour as the Group By selection and want the top 25 groups for a 24-hr time period, the results of the report will be generated on an hourly basis over a 24-hr period. The first column in the report will be the hour and the next set of columns will be the rest of your selected report columns.
	The following example illustrates how the **Selected Columns** and **Sort By**/**Group By** criteria work together when generating reports:

© 2018 Palo Alto Networks, Inc.

Selection	Description
	 The columns circled in red (above) depict the columns selected, which are the attributes that you match against for generating the report. Each log entry from the data source is parsed and these columns are matched on. If multiple sessions have the same values for the selected columns, the sessions are aggregated and the repeat count (or sessions) is incremented. The column circled in blue indicates the chosen sort order. When the sort order (**Sort By**) is specified, the data is sorted (and aggregated) by the selected attribute. The column circled in green indicates the **Group By** selection, which serves as an anchor for the report. The **Group By** column is used as a match criteria to filter for the top N groups. Then, for each of the top N groups, the report enumerates the values for all the other selected columns.
	For example, if a report has the following selections:  The output will display as follows:  The report is anchored by **Day** and sorted by **Sessions**. It lists the 5 days (**5 Groups**) with maximum traffic in the **Last 7 Days** time frame. The data is enumerated by the **Top 5** sessions for each day for the selected columns—**App Category**, **App Subcategory** and **Risk**.
Time Frame	The date range for which you want to analyze data. You can define a custom range or select a time period ranging from the last 15 minutes to the last 30

© 2018 Palo Alto Networks, Inc.

Selection	Description
	days. The reports can be run on demand or scheduled to run at a daily or weekly cadence.
Query Builder	The query builder allows you to define specific queries to further refine the selected attributes. It allows you see just what you want in your report using **and** and **or** operators and a match criteria, and then include or exclude data that matches or negates the query in the report. Queries enable you to generate a more focused collation of information in a report.

Generate Custom Reports

You can configure custom reports that the firewall generates immediately (on demand) or on schedule (each night). To understand the selections available to create a purposeful custom report, see Custom Reports.



After the firewall has generated a scheduled custom report, you risk invalidating the past results of that report if you modify its configuration to change its future output. If you need to modify a scheduled report configuration, the best practice is to create a new report.

STEP 1 | Select **Monitor** > **Manage Custom Reports**.

STEP 2 | Click **Add** and then enter a **Name** for the report.



To base a report on an predefined template, click Load Template and choose the template. You can then edit the template and save it as a custom report.

STEP 3 | Select the **Database** to use for the report.



Each time you create a custom report, a log view report is automatically created. This report show the logs that were used to build the custom report. The log view report uses the same name as the custom report, but appends the phrase (Log View) to the report name.

When creating a report group, you can include the log view report with the custom report. For more information, see Manage Report Groups.

STEP 4 | Select the **Scheduled** check box to run the report each night. The report is then available for viewing in the **Reports** column on the side.

STEP 5 | Define the filtering criteria. Select the **Time Frame**, the **Sort By** order, **Group By** preference, and select the columns that must display in the report.

STEP 6 | **(**Optional**)** Select the **Query Builder** attributes if you want to further refine the selection criteria. To build a report query, specify the following and click **Add**. Repeat as needed to construct the full query.

- **Connector**—Choose the connector (and/or) to precede the expression you are adding.
- **Negate**—Select the check box to interpret the query as a negation. If, for example, you choose to match entries in the last 24 hours and/or are originating from the untrust zone, the negate option causes a match on entries that are not in the past 24 hours and/or are not from the untrust zone.
- **Attribute**—Choose a data element. The available options depend on the choice of database.

© 2018 Palo Alto Networks, Inc.

- **Operator**—Choose the criterion to determine whether the attribute applies (such as =). The available options depend on the choice of database.
- **Value**—Specify the attribute value to match.

For example, the following figure (based on the `Traffic Log` database) shows a query that matches if the Traffic log entry was received in the past 24 hours and is from the untrust zone.

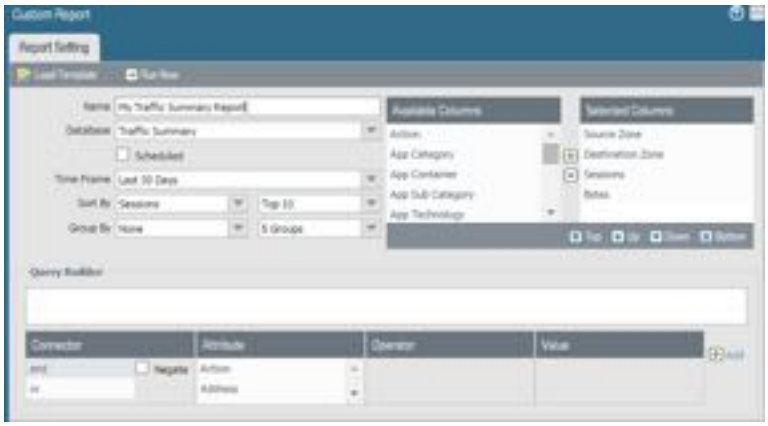


STEP 7 | To test the report settings, select **Run Now**. Modify the settings as required to change the information that is displayed in the report.

STEP 8 | Click **OK** to save the custom report.

Examples of Custom Reports

If you want to set up a simple report in which you use the traffic summary database from the last 30 days, and sort the data by the top 10 sessions and these sessions are grouped into 5 groups by day of the week. You would set up the custom report to look like this:



And the PDF output for the report would look as follows:

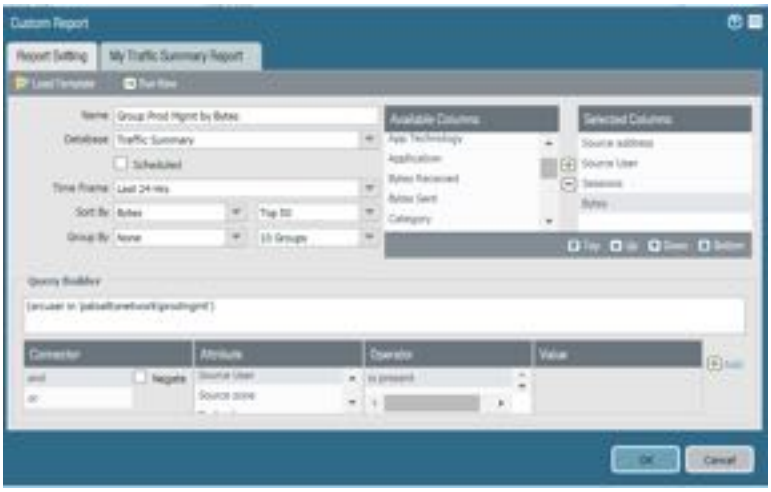
My Traffic Summary Report

ca1demo.paloaltonetworks.com : 2016/01/25 10:34:39 - 2016/02/24 10:34:38

Source Zone	Destination Zone	App Category	Application	Sessions	Bytes
Tap	Tap	general-internet	web-browsing	74.54 M	2.47 T
Tap	Tap	networking	dns	52.03 M	28.93 G
Tap	Tap	networking	ssl	18.01 M	678.13 G
Tap	Tap	general-internet	bittorrent	9.80 M	1.62 T
Tap	Tap	general-internet	google-base	4.48 M	168.99 G
Tap	Tap	unknown	insufficient-data	4.45 M	31.30 G
Tap	Tap	collaboration	facebook-base	4.09 M	99.14 G
Tap	Tap	networking	ntp	4.07 M	3.29 G
Tap	Tap	collaboration	blackboard	2.84 M	186 G
Tap	Tap	collaboration	smtp	1.92 M	172.57 G
Tap	Tap	networking	icmp	1.36 M	320.49 M
Tap	Tap	general-internet	gnutella	1.17 M	17.84 G
Tap	Tap	collaboration	myspace-base	1.10 M	35.22 G
Tap	Tap	general-internet	ping	1.06 M	86.21 M
Tap	Tap	general-internet	flash	1.01 M	168.14 G

© 2018 Palo Alto Networks, Inc.

Now, if you want to use the query builder to generate a custom report that represents the top consumers of network resources within a user group, you would set up the report to look like this:



The report would display the top users in the product management user group sorted by bytes.

Generate Botnet Reports

The botnet report enables you to use heuristic and behavior-based mechanisms to identify potential malware- or botnet-infected hosts in your network. To evaluate botnet activity and infected hosts, the firewall correlates user and network activity data in Threat, URL, and Data Filtering logs with the list of malware URLs in PAN-DB, known dynamic DNS domain providers, and domains registered within the last 30 days. You can configure the report to identify hosts that visited those sites, as well as hosts that communicated with Internet Relay Chat (IRC) servers or that used unknown applications. Malware often use dynamic DNS to avoid IP blacklisting, while IRC servers often use bots for automated functions.

The firewall requires Threat Prevention and URL Filtering licenses to use the botnet report. You can Use the Automated Correlation Engine *on page 334 to monitor suspicious activities based on additional indicators besides those that the botnet report uses. However, the botnet report is the only tool that uses newly registered domains as an indicator.*

- Configure a Botnet Report on page 367
- Interpret Botnet Report Output on page 368

Configure a Botnet Report

You can schedule a botnet report or run it on demand. The firewall generates scheduled botnet reports every 24 hours because behavior-based detection requires correlating traffic across multiple logs over that timeframe.

STEP 1 | Define the types of traffic that indicate possible botnet activity.

1. Select **Monitor** > **Botnet** and click **Configuration** on the right side of the page.
2. **Enable** and define the **Count** for each type of HTTP Traffic that the report will include.

   The **Count** values represent the minimum number of events of each traffic type that must occur for the report to list the associated host with a higher confidence score (higher likelihood of botnet infection). If the number of events is less than the **Count**, the report will display a lower confidence score or (for certain traffic types) won't display an entry for the host. For example, if you set the **Count** to three for **Malware URL visit**, then hosts that visit three or more known malware URLs

© 2018 Palo Alto Networks, Inc.

will have higher scores than hosts that visit less than three. For details, see Interpret Botnet Report Output on page 368.

3. Define the thresholds that determine whether the report will include hosts associated with traffic involving Unknown TCP or Unknown UDP applications.
4. Select the **IRC** check box to include traffic involving IRC servers.
5. Click **OK** to save the report configuration.

STEP 2 | Schedule the report or run it on demand.

1. Click **Report Setting** on the right side of the page.
2. Select a time interval for the report in the **Test Run Time Frame** drop-down.
3. Select the **No. of Rows** to include in the report.
4. (Optional) **Add** queries to the Query Builder to filter the report output by attributes such as source/destination IP addresses, users, or zones.

   For example, if you know in advance that traffic initiated from the IP address 10.3.3.15 contains no potential botnet activity, add **`not (addr.src in 10.0.1.35)`** as a query to exclude that host from the report output. For details, see Interpret Botnet Report Output on page 368.
5. Select **Scheduled** to run the report daily or click **Run Now** to run the report immediately.
6. Click **OK** and **Commit**.

Interpret Botnet Report Output

The botnet report displays a line for each host that is associated with traffic you defined as suspicious when configuring the report. For each host, the report displays a confidence score of 1 to 5 to indicate the likelihood of botnet infection, where 5 indicates the highest likelihood. The scores correspond to threat severity levels: 1 is informational, 2 is low, 3 is medium, 4 is high, and 5 is critical. The firewall bases the scores on:

- **Traffic type**—Certain HTTP traffic types are more likely to involve botnet activity. For example, the report assigns a higher confidence to hosts that visit known malware URLs than to hosts that browse to IP domains instead of URLs, assuming you defined both those activities as suspicious.
- **Number of events**—Hosts that are associated with a higher number of suspicious events will have higher confidence scores based on the thresholds (**Count** values) you define when you Configure a Botnet Report on page 367.
- **Executable downloads**—The report assigns a higher confidence to hosts that download executable files. Executable files are a part of many infections and, when combined with the other types of suspicious traffic, can help you prioritize your investigations of compromised hosts.

When reviewing the report output, you might find that the sources the firewall uses to evaluate botnet activity (for example, the list of malware URLs in PAN-DB) have gaps. You might also find that these sources identify traffic that you consider safe. To compensate in both cases, you can add query filters when you Configure a Botnet Report on page 367.

Generate the SaaS Application Usage Report

The SaaS Application Usage PDF report is a two-part report that is based on the notion of sanctioned and unsanctioned applications. A sanctioned application is an application that you formally approve for use on your network; a SaaS application is an application that has the characteristic SaaS=yes in the applications details page in **Objects** > **Applications**, all other applications are considered as non-SaaS. To indicate that you have sanctioned a SaaS or non-SaaS application, you must tag it with the new predefined tag named Sanctioned. The firewall and Panorama consider any application without this predefined tag as unsanctioned for use on the network.

© 2018 Palo Alto Networks, Inc.

- The first part of the report (10 pages) focuses on the SaaS applications used on your network during the reporting period. It presents a comparison of sanctioned versus unsanctioned SaaS applications by total number of applications used on your network, bandwidth consumed by these applications, the number of users using these applications, top user groups that use the largest number of SaaS applications, and the top user groups that transfer the largest volume of data through sanctioned and unsanctioned SaaS applications. This first part of the report also highlights the top SaaS application subcategories listed in order by maximum number of applications used, the number of users, and the amount of data (bytes) transferred in each application subcategory.
- The second part of the report focuses on the detailed browsing information for SaaS and non-SaaS applications for each application subcategory listed in the first-part of the report. For each application in a subcategory, it also includes information about the top users who transferred data, the top blocked or alerted file types, and the top threats for each application. In addition, this section of the report tallies samples for each application that the firewall submitted for WildFire analysis, and the number of samples determined to be benign and malicious.

Use the insights from this report to consolidate the list of business-critical and approved SaaS applications and to enforce policies for controlling unsanctioned applications that pose an unnecessary risk for malware propagation and data leaks.



The predefined SaaS application usage report introduced in PAN-OS 7.0 is still available as a daily report that lists the top 100 SaaS applications (with the SaaS application characteristic, SaaS=yes) running on your network on a given day.

STEP 1 | Tag applications that you approve for use on your network as Sanctioned.



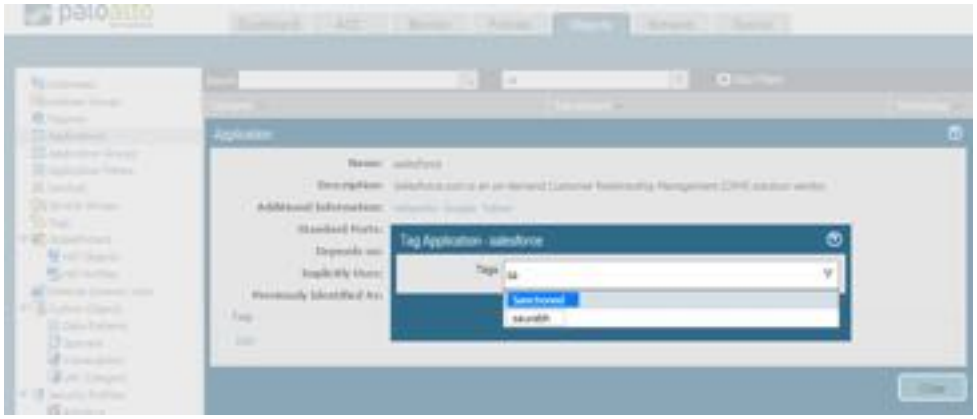
For generating an accurate and informative report, you need to tag the sanctioned applications consistently across firewalls with multiple virtual systems, and across firewalls that belong to a device group on Panorama. If the same application is tagged as sanctioned in one virtual system and is not sanctioned in another or, on Panorama, if an application is unsanctioned in a parent device group but is tagged as sanctioned in a child device group (or vice versa), the SaaS Application Usage report will report the application as partially sanctioned and will have overlapping results.

Example: If Box is sanctioned on vsys1 and Google Drive is sanctioned on vsys2, Google Drive users in vsys1 will be counted as users of an unsanctioned SaaS application and Box users in vsys2 will be counted as users of an unsanctioned SaaS application. The key finding in the report will highlight that a total of two unique SaaS applications are discovered on the network with two sanctioned applications and two unsanctioned applications.

1. Select **Object** > **Applications**.
2. Click the application **Name** to edit an application and select **Edit** in the Tag section.
3. Select **Sanctioned** from the **Tags** drop-down.

   You must use the predefined **Sanctioned** tag (with the green colored background). If you use any other tag to indicate that you sanctioned an application, the firewall will fail to recognize the tag and the report will be inaccurate.

© 2018 Palo Alto Networks, Inc.



4. Click **OK** and **Close** to exit all open dialogs.

STEP 2 | Configure the SaaS Application Usage report.

1. Select **Monitor** > **PDF Reports** > **SaaS Application Usage**.
2. Click **Add**, enter a **Name**, and select a **Time Period** for the report (default is **Last 7 Days**).

   By default, the report includes detailed information on the top SaaS and non-SaaS application subcategories, which can make the report large by page count and file size. Clear the Include detailed application category information in report check box if you want to reduce the file size and restrict the page count to 10 pages.
3. Select whether you want the report to **Include logs from**:
   - **All User Groups and Zones**—The report includes data on all security zones and user groups available in the logs.

     If you want to include specific user groups in the report, select **Include user group information in the report** and click the **manage groups** link to select the groups you want to include. You must add between one and up to a maximum of 25 user groups, so that the firewall or Panorama can filter the logs for the selected user groups. If you do select the groups to include, the report will aggregate all user groups in to one group called Others.
   - **Selected Zone**—The report filters data for the specified security zone, and includes data on that zone only.

     If you want to include specific user groups in the report, select **Include user group information in the report** and click the **manage groups for selected zone** link to select the user groups within this zone that you want to include in the report. You must add between one and up to a maximum of 25 user groups, so that the firewall or Panorama can filter the logs for the selected user groups within the security zone. If you do select the groups to include, the report will aggregate all user groups in to one group called Others.
   - **Selected User Group**—The report filters data for the specified user group only, and includes SaaS application usage information for the selected user group only.

© 2018 Palo Alto Networks, Inc.



4. Select whether you want to include all the application subcategories in the report (the default) or **Limit the max subcategories in the report** to the top 10, 15, 20 or 25 categories (default is all subcategories).
5. Click **Run Now** to generate the report on-demand for the last 7-day and the last 30-day time period. Make sure that the pop-up blocker is disabled on your browser because the report opens in a new tab.
6. Click **OK** to save your changes.

STEP 3 | Schedule Reports for Email Delivery on page 375.

The last 90-days report must be scheduled for email delivery.

On the PA-200, PA-220, and PA-500 firewalls, the SaaS Application Usage report is not sent as a PDF attachment in the email. Instead, the email includes a link that you must click to open the report in a web browser.

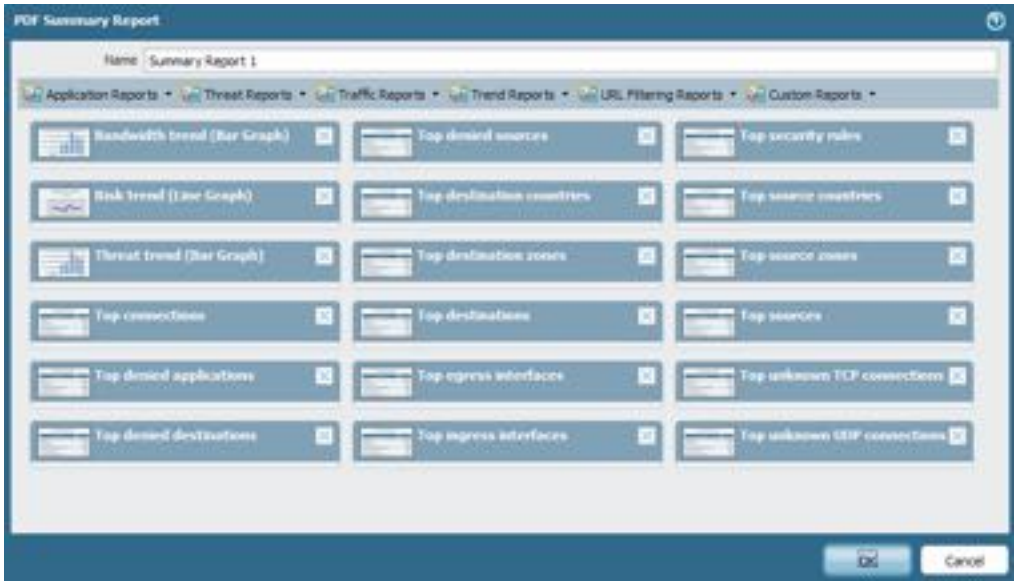
Manage PDF Summary Reports

PDF summary reports contain information compiled from existing reports, based on data for the top 5 in each category (instead of top 50). They also contain trend charts that are not available in other reports.

STEP 1 | Set up a **PDF Summary Report**.

1. Select **Monitor** > **PDF Reports** > **Manage PDF Summary**.
2. Click **Add** and then enter a **Name** for the report.
3. Use the drop-down for each report group and select one or more of the elements to design the PDF Summary Report. You can include a maximum of 18 report elements.

© 2018 Palo Alto Networks, Inc.



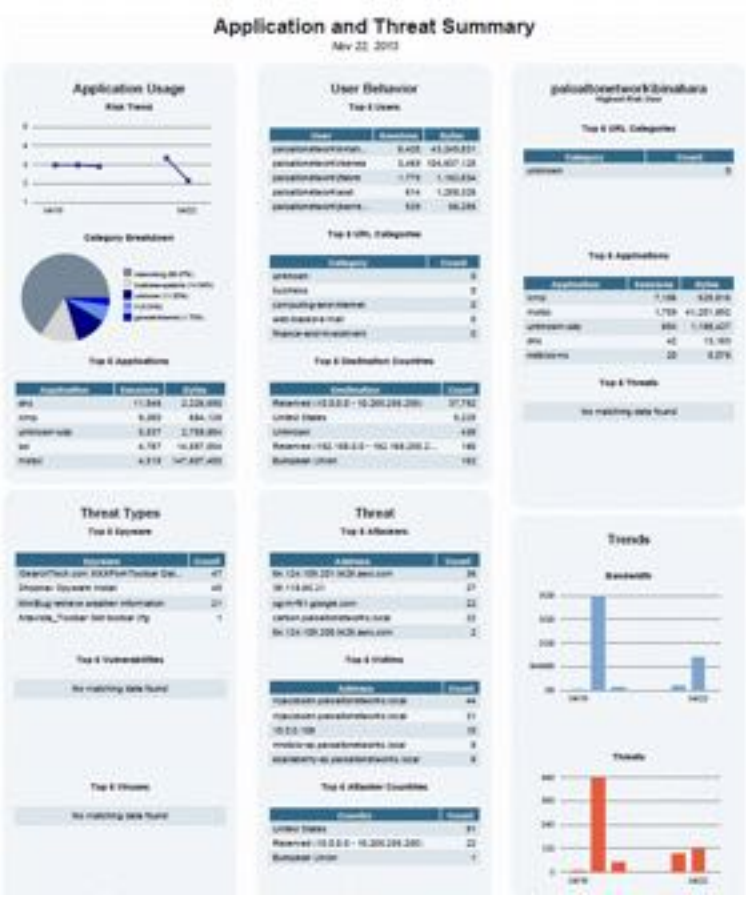
- To remove an element from the report, click the **x** icon or clear the selection from the drop-down for the appropriate report group.
- To rearrange the reports, drag and drop the element icons to another area of the report.

4. Click **OK t**o save the report.
5. **Commit** the changes.

STEP 2 | View the report.

To download and view the PDF Summary Report, see View Reports on page 361.

© 2018 Palo Alto Networks, Inc.



Generate User/Group Activity Reports

User/Group Activity reports summarize the web activity of individual users or user groups. Both reports include the same information except for the **Browsing Summary by URL Category** and **Browse time calculations**, which only the User Activity report includes.

You must configure User-ID on page 445 on the firewall to access the list of users and user groups.

STEP 1 | Configure the browse times and number of logs for User/Group Activity reports.

Required only if you want to change the default values.

1. Select **Device** > **Setup** > **Management**, edit the Logging and Reporting Settings, and select the **Log Export and Reporting** tab.
2. For the **Max Rows in User Activity Report**, enter the maximum number of rows that the detailed user activity report supports (range is 1-1048576, default is 5000). This determines the number of logs that the report analyzes.
3. Enter the **Average Browse Time** in seconds that you estimate users should take to browse a web page (range is 0-300, default is 60). Any request made after the average browse time elapses is considered a new browsing activity. The calculation uses Container Pages on page 651 (logged in the URL Filtering logs) as the basis and ignores any new web pages that are loaded between the time of the first request (start time) and the average browse time. For example, if you set the **Average Browse Time** to two minutes and a user opens a web page and views that page for five minutes, the browse time for that page will still be two minutes. This is done because the firewall can't determine how long a user views a given page. The average browse time calculation ignores sites categorized as web advertisements and content delivery networks.

© 2018 Palo Alto Networks, Inc.

4. For the **Page Load Threshold**, enter the estimated time in seconds for page elements to load on the page (default is 20). Any requests that occur between the first page load and the page load threshold are assumed to be elements of the page. Any requests that occur outside of the page load threshold are assumed to be the user clicking a link within the page.
5. Click **OK** to save your changes.

STEP 2 | Generate the User/Group Activity report.

1. Select **Monitor** > **PDF Reports** > **User Activity Report**.
2. Click **Add** and then enter a **Name** for the report.
3. Create the report:
   - User Activity Report—Select **User** and enter the **Username** or **IP address** (IPv4 or IPv6) of the user.
   - Group Activity Report—Select **Group** and select the **Group Name** of the user group.
4. Select the **Time Period** for the report.
5. (Optional) Select the **Include Detailed Browsing** check box (default is cleared) to include detailed URL logs in the report.

   The detailed browsing information can include a large volume of logs (thousands of logs) for the selected user or user group and can make the report very large.
6. To run the report on demand, click **Run Now**.
7. To save the report configuration, click **OK**. You can't save the output of User/Group Activity reports on the firewall. To schedule the report for email delivery, see Schedule Reports for Email Delivery on page 375.

Manage Report Groups

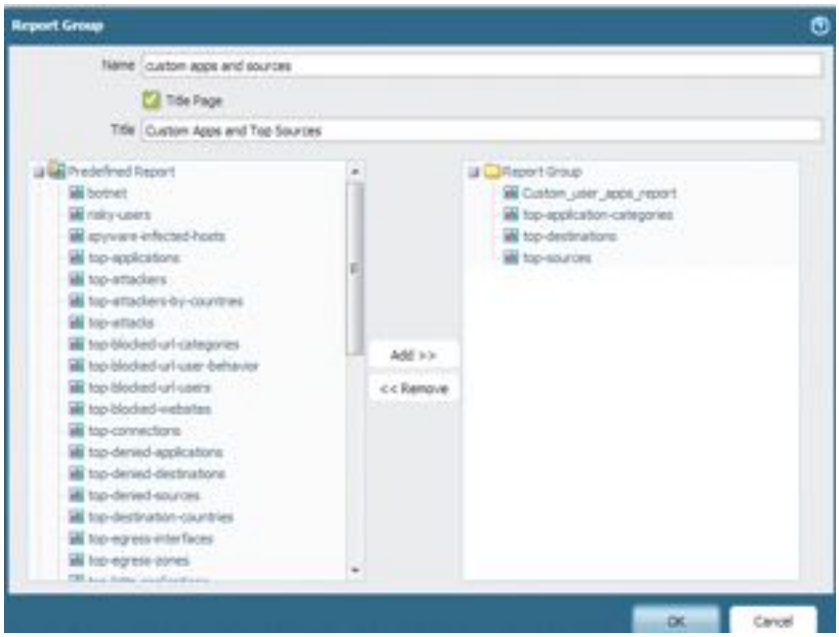
Report groups allow you to create sets of reports that the system can compile and send as a single aggregate PDF report with an optional title page and all the constituent reports included.

Set up report groups.

You must set up a **Report Group** to email report(s).

1. Create an Email server profile.
2. Define the **Report Group**. A report group can compile predefined reports, PDF Summary reports, custom reports, and Log View report into a single PDF.
   1. Select **Monitor** > **Report Group**.
   2. Click **Add** and then enter a **Name** for the report group.
   3. (Optional) Select **Title Page** and add a **Title** for the PDF output.
   4. Select reports from the left column and click **Add** to move each report to the report group on the right.

© 2018 Palo Alto Networks, Inc.



The **Log View** report is a report type that is automatically created each time you create a custom report and uses the same name as the custom report. This report will show the logs that were used to build the contents of the custom report.

To include the log view data, when creating a report group, add your custom report under the **Custom Reports** list and then add the log view report by selecting the matching report name from the **Log View** list. The report will include the custom report data and the log data that was used to create the custom report.

5. Click **OK** to save the settings.
6. To use the report group, see Schedule Reports for Email Delivery.

Schedule Reports for Email Delivery

Reports can be scheduled for daily delivery or delivered weekly on a specified day. Scheduled reports are executed starting at 2:00 AM, and email delivery starts after all scheduled reports have been generated.

STEP 1 | Select **Monitor** > **PDF Reports** > **Email Scheduler** and click **Add**.

STEP 2 | Enter a **Name** to identify the schedule.

STEP 3 | Select the **Report Group** for email delivery. To set up a report group; see Manage Report Groups.

STEP 4 | For the **Email Profile**, select an Email server profile to use for delivering the reports, or click the **Email Profile** link to Create an Email server profile.

STEP 5 | Select the frequency at which to generate and send the report in **Recurrence**.

STEP 6 | The **Override Email Addresses** field allows you to send this report exclusively to the specified recipients. When you add recipients to the field, the firewall does not send the report to the recipients configured in the Email server profile. Use this option for those occasions when the report is for the attention of someone other than the administrators or recipients defined in the Email server profile.

© 2018 Palo Alto Networks, Inc.

STEP 7 | Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Use External Services for Monitoring

Using an external service to monitor the firewall enables you to receive alerts for important events, archive monitored information on systems with dedicated long-term storage, and integrate with third-party security monitoring tools. The following are some common scenarios for using external services:

- For immediate notification about important system events or threats, you can Monitor Statistics Using SNMP on page 422, Forward Traps to an SNMP Manager on page 423, or Configure Email Alerts on page 381.
- To send an HTTP-based API request directly to any third-party service that exposes an API to automate a workflow or an action. You can, for example, forward logs that match a defined criteria to create an incidence ticket on Service Now instead of relying on an external system to convert syslog messages or SNMP traps to an HTTP request. You can modify the URL, HTTP header, parameters, and the payload in the HTTP request to trigger an action based on the attributes in a firewall log. See Forward Logs to an HTTP(S) Destination on page 433.
- For long-term log storage and centralized firewall monitoring, you can Configure Syslog Monitoring on page 382 to send log data to a syslog server. This enables integration with third-party security monitoring tools such as Splunk! or ArcSight.
- For monitoring statistics on the IP traffic that traverses firewall interfaces, you can Configure NetFlow Exports on page 436 to view the statistics in a NetFlow collector.

You can Configure Log Forwarding on page 378 from the firewalls directly to external services or from the firewalls to Panorama and then configure Panorama to forward logs to the servers. Refer to Log Forwarding Options for the factors to consider when deciding where to forward logs.



You can't aggregate NetFlow records on Panorama; you must send them directly from the firewalls to a NetFlow collector.

© 2018 Palo Alto Networks, Inc.

Configure Log Forwarding

In an environment where you use multiple firewalls to control and analyze network traffic, any single firewall can display logs and reports only for the traffic it monitors. Because logging in to multiple firewalls can make monitoring a cumbersome task, you can more efficiently achieve global visibility into network activity by forwarding the logs from all firewalls to Panorama or external services. If you Use External Services for Monitoring, the firewall automatically converts the logs to the necessary format: syslog messages, SNMP traps, email notifications, or as an HTTP payload to send the log details to an HTTP(S) server. In cases where some teams in your organization can achieve greater efficiency by monitoring only the logs that are relevant to their operations, you can create forwarding filters based on any log attributes (such as threat type or source user). For example, a security operations analyst who investigates malware attacks might be interested only in Threat logs with the type attribute set to wildfire-virus.



You can forward logs from the firewalls directly to external services or from the firewalls to Panorama and then configure Panorama to forward logs to the servers. Refer to Log Forwarding Options for the factors to consider when deciding where to forward logs.

You can use Secure Copy (SCP) commands from the CLI to export the entire log database to an SCP server and import it to another firewall. Because the log database is too large for an export or import to be practical on the PA-7000 Series firewall, it does not support these options. You can also use the web interface on all platforms to View and Manage Reports, but only on a per log type basis, not for the entire log database.

STEP 1 | Configure a server profile for each external service that will receive log information.



You can use separate profiles to send different sets of logs, filtered by log attributes, to a different server. To increase availability, define multiple servers in a single profile.

Configure one or more of the following server profiles:

- Create an Email server profile.
- Step Configure an SNMP Trap server profile. To enable the SNMP manager (trap server) to interpret firewall traps, you must load the Palo Alto Networks Supported MIBs into the SNMP manager and, if necessary, compile them. For details, refer to your SNMP management software documentation.
- If the syslog server requires client authentication, you must also Create a certificate to secure syslog communication over TLSv1.2.
- Configure an HTTP server profile (see Forward Logs to an HTTP(S) Destination).

STEP 2 | Create a Log Forwarding profile.

The profile defines the destinations for Traffic, Threat, WildFire Submission, URL Filtering, Data Filtering, Tunnel and Authentication logs.

1. Select **Objects** > **Log Forwarding** and **Add** a profile.
2. Enter a **Name** to identify the profile.

   If you want the firewall to automatically assign the profile to new security rules and zones, enter **`default`**. If you don't want a default profile, or you want to override an existing default profile, enter a **Name** that will help you identify the profile when assigning it to security rules and zones.

© 2018 Palo Alto Networks, Inc.

If no log forwarding profile named ***default*** *exists, the profile selection is set to None by default in new security rules (Log Forwarding field) and new security zones (Log Setting field), although you can change the selection.*

3. **Add** one or more *match list profiles*.

   The profiles specify log query filters, forwarding destinations, and automatic actions such as tagging. For each match list profile:

   1. Enter a **Name** to identify the profile.
   2. Select the **Log Type**.
   3. In the **Filter** drop-down, select **Filter Builder**. Specify the following and then **Add** each query:
      - **Connector** logic (and/or)
      - Log **Attribute**
      - **Operator** to define inclusion or exclusion logic
      - Attribute **Value** for the query to match
   4. Select Panorama if you want to forward logs to Log Collectors or the Panorama management server.
   5. For each type of external service that you use for monitoring (SNMP, Email, Syslog, and HTTP), **Add** one or more server profiles.
4. Click **OK** to save the Log Forwarding profile.

STEP 3 | Assign the Log Forwarding profile to policy rules and network zones.

Security, Authentication, and DoS Protection rules support log forwarding. In this example, you assign the profile to a Security rule.

Perform the following steps for each rule that you want to trigger log forwarding:

1. Select **Policies** > **Security** and edit the rule.
2. Select **Actions** and select the **Log Forwarding** profile you created.
3. Set the **Profile Type** to **Profiles** or **Group**, and then select the security profiles or **Group Profile** required to trigger log generation and forwarding for:
   - Threat logs—Traffic must match any security profile assigned to the rule.
   - WildFire Submission logs—Traffic must match a WildFire Analysis profile assigned to the rule.
4. For Traffic logs, select **Log At Session Start** and/or **Log At Session End**.
5. Click **OK** to save the rule.

STEP 4 | Configure the destinations for System, Configuration, User-ID, HIP Match, and Correlation logs.

Panorama generates Correlation logs based on the firewall logs it receives, rather than aggregating Correlation logs from firewalls.

1. Select **Device** > **Log Settings**.
2. For each log type that the firewall will forward, see Step Add one or more match list profiles.

STEP 5 | (PA-7000 Series firewalls only) Configure a log card interface to perform log forwarding.

1. Select **Network** > **Interfaces** > **Ethernet** and click **Add Interface**.
2. Select the **Slot** and **Interface Name**.
3. Set the **Interface Type** to **Log Card**.
4. Enter the **IP Address**, **Default Gateway**, and (for IPv4 only) **Netmask**.
5. Select **Advanced** and specify the **Link Speed**, **Link Duplex**, and **Link State**.

© 2018 Palo Alto Networks, Inc.



These fields default to auto, which specifies that the firewall automatically determines the values based on the connection. However, the minimum recommended Link Speed for any connection is 1000 (Mbps).

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

STEP 6 | Commit and verify your changes.

1. **Commit** your changes.
2. Verify the log destinations you configured are receiving firewall logs:
   - Panorama—If the firewall forwards logs to a Panorama virtual appliance in Panorama mode or to an M-Series appliance, you must configure a Collector Group before Panorama will receive the logs. You can then verify log forwarding.
   - Email server—Verify that the specified recipients are receiving logs as email notifications.
   - Syslog server—Refer to your syslog server documentation to verify it's receiving logs as syslog messages.
   - SNMP manager—Use an SNMP Manager to Explore MIBs and Objects to verify it's receiving logs as SNMP traps.
   - HTTP server—Forward Logs to an HTTP(S) Destination.

© 2018 Palo Alto Networks, Inc.

Configure Email Alerts

You can configure email alerts for System, Config, HIP Match, Correlation, Threat, WildFire Submission, and Traffic logs.

STEP 1 | Create an Email server profile.



You can use separate profiles to send email notifications for each log type to a different server. To increase availability, define multiple servers (up to four) in a single profile.

1. Select **Device** > **Server Profiles** > **Email**.
2. Click **Add** and then enter a **Name** for the profile.
3. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available.
4. For each Simple Mail Transport Protocol (SMTP) server (email server), click **Add** and define the following information:
   - **Name**—Name to identify the SMTP server (1-31 characters). This field is just a label and doesn't have to be the hostname of an existing email server.
   - **Email Display Name**—The name to show in the From field of the email.
   - **From**—The email address from which the firewall sends emails.
   - **To**—The email address to which the firewall sends emails.
   - **Additional Recipient**—If you want to send emails to a second account, enter the address here. You can add only one additional recipient. For multiple recipients, add the email address of a distribution list.
   - **Email Gateway**—The IP address or hostname of the SMTP gateway to use for sending emails.
5. (Optional) Select the **Custom Log Format** tab and customize the format of the email messages. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide.
6. Click **OK** to save the Email server profile.

STEP 2 | Configure email alerts for Traffic, Threat, and WildFire Submission logs.

1. See Step Create a Log Forwarding profile.
   1. Select **Objects** > **Log Forwarding**, click **Add**, and enter a **Name** to identify the profile.
   2. For each log type and each severity level or WildFire verdict, select the Email server profile and click **OK**.
2. See Step Assign the Log Forwarding profile to policy rules and network zones.

STEP 3 | Configure email alerts for System, Config, HIP Match, and Correlation logs.

1. Select **Device** > **Log Settings**.
2. For System and Correlation logs, click each Severity level, select the **Email** server profile, and click **OK**.
3. For Config and HIP Match logs, edit the section, select the **Email** server profile, and click **OK**.
4. Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Use Syslog for Monitoring

Syslog is a standard log transport mechanism that enables the aggregation of log data from different network devices—such as routers, firewalls, printers—from different vendors into a central repository for archiving, analysis, and reporting. Palo Alto Networks firewalls can forward every type of log they generate to an external syslog server. You can use TCP or TLS (TLSv1.2 only) for reliable and secure log forwarding, or UDP for non-secure forwarding.

- Configure Syslog Monitoring
- Syslog Field Descriptions

Configure Syslog Monitoring

To Use Syslog for Monitoring a Palo Alto Networks firewall, create a Syslog server profile and assign it to the log settings for each log type. Optionally, you can configure the header format used in syslog messages and enable client authentication for syslog over TLSv1.2.

STEP 1 | Configure a Syslog server profile.



You can use separate profiles to send syslogs for each log type to a different server. To increase availability, define multiple servers (up to four) in a single profile.

1. Select **Device** > **Server Profiles** > **Syslog**.
2. Click **Add** and enter a **Name** for the profile.
3. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available.
4. For each syslog server, click **Add** and enter the information that the firewall requires to connect to it:
   - **Name**—Unique name for the server profile.
   - **Syslog Server**—IP address or fully qualified domain name (FQDN) of the syslog server.
   - **Transport**—Select **TCP**, **UDP**, or **SSL** (TLS) as the protocol for communicating with the syslog server. For **SSL**, the firewall supports only TLSv1.2.
   - **Port**—The port number on which to send syslog messages (default is UDP on port 514); you must use the same port number on the firewall and the syslog server.
   - **Format**—Select the syslog message format to use: **BSD** (the default) or **IETF**. Traditionally, **BSD** format is over UDP and **IETF** format is over TCP or SSL/TLS.
   - **Facility**—Select a syslog standard value (default is **LOG_USER**) to calculate the priority (PRI) field in your syslog server implementation. Select the value that maps to how you use the PRI field to manage your syslog messages.
5. (**Optional**) To customize the format of the syslog messages that the firewall sends, select the **Custom Log Format** tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide.
6. Click **OK** to save the server profile.

STEP 2 | Configure syslog forwarding for Traffic, Threat, and WildFire Submission logs.

1. See Step Create a Log Forwarding profile.
   1. Select **Objects** > **Log Forwarding**, click **Add**, and enter a **Name** to identify the profile.
   2. For each log type and each severity level or WildFire verdict, select the **Syslog** server profile and click **OK**.
2. See Step Assign the Log Forwarding profile to policy rules and network zones.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Configure syslog forwarding for System, Config, HIP Match, and Correlation logs.

1. Select **Device** > **Log Settings**.
2. For System and Correlation logs, click each Severity level, select the **Syslog** server profile, and click **OK**.
3. For Config, HIP Match, and Correlation logs, edit the section, select the **Syslog** server profile, and click **OK**.

STEP 4 | (**Optional**) Configure the header format of syslog messages.

The log data includes the unique identifier of the firewall that generated the log. Choosing the header format provides more flexibility in filtering and reporting on the log data for some Security Information and Event Management (SIEM) servers.

This is a global setting and applies to all Syslog server profiles configured on the firewall.

1. Select **Device** > **Setup** > **Management** and edit the Logging and Reporting Settings.
2. Select the **Log Export and Reporting** tab and select the Syslog HOSTNAME Format:
   - **FQDN** (default)—Concatenates the hostname and domain name defined on the sending firewall.
   - **hostname**—Uses the hostname defined on the sending firewall.
   - **ipv4-address**—Uses the IPv4 address of the firewall interface used to send logs. By default, this is the MGT interface.
   - **ipv6-address**—Uses the IPv6 address of the firewall interface used to send logs. By default, this is the MGT interface.
   - **none**—Leaves the hostname field unconfigured on the firewall. There is no identifier for the firewall that sent the logs.
3. Click **OK** to save your changes.

STEP 5 | Create a certificate to secure syslog communication over TLSv1.2.

Required only if the syslog server uses client authentication. The syslog server uses the certificate to verify that the firewall is authorized to communicate with the syslog server.

Ensure the following conditions are met:

- The private key must be available on the sending firewall; the keys can't reside on a Hardware Security Module (HSM).
- The subject and the issuer for the certificate must not be identical.
- The syslog server and the sending firewall must have certificates that the same trusted certificate authority (CA) signed. Alternatively, you can generate a self-signed certificate on the firewall, export the certificate from the firewall, and import it in to the syslog server.
- Beginning in PAN-OS 8.0.6, the connection to a Syslog server over TLS is validated using the Online Certificate Status Protocol (OCSP) when available. To prevent certificate verification failures, ensure that the certificate chain is valid and can be verified using OCSP. You cannot bypass OCSP failures for secure connections.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Generate**.
2. Enter a **Name** for the certificate.
3. In the **Common Name** field, enter the IP address of the firewall sending logs to the syslog server.
4. In **Signed by**, select the trusted CA or the self-signed CA that the syslog server and the sending firewall both trust.

   The certificate can't be a **Certificate Authority** nor an **External Authority** (certificate signing request [CSR]).
5. Click **Generate**. The firewall generates the certificate and key pair.
6. Click the certificate Name to edit it, select the **Certificate for Secure Syslog** check box, and click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 6 | Commit your changes and review the logs on the syslog server.

1. Click **Commit**.
2. To review the logs, refer to the documentation of your syslog management software. You can also review the Syslog Field Descriptions.

Syslog Field Descriptions

The following topics list the standard fields of each log type that Palo Alto Networks firewalls can forward to an external server, as well as the severity levels, custom formats, and escape sequences. To facilitate parsing, the delimiter is a comma: each field is a comma-separated value (CSV) string. The FUTURE_USE tag applies to fields that the firewalls do not currently implement.

WildFire Submissions logs are a subtype of Threat log and use the same syslog format.

- Traffic Log Fields
- Threat Log Fields
- HIP Match Log Fields
- User-ID Log Fields
- Tunnel Inspection Log Fields
- Config Log Fields
- Authentication Log Fields
- System Log Fields
- GTP Log Fields
- Correlated Events Log Fields
- Custom Log/Event Format
- Escape Sequences

Traffic Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Bytes, Bytes Sent, Bytes Received, Packets, Start Time, Elapsed Time, Category, FUTURE_USE, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Packets Sent, Packets Received, Session End Reason, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Action Source, Source VM UUID, Destination VM UUID, Tunnel ID/IMSI, Monitor Tag/IMEI, Parent Session ID, Parent Start Time, Tunnel Type

Field Name	Description
Receive Time	Time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the firewall that generated the log.
Type	Specifies type of log; values are traffic, threat, config, system and hip-match.
Threat/Content Type	Subtype of traffic log; values are start, end, drop, and deny

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	• Start—session started • End—session ended • Drop—session dropped before the application is identified and there is no rule that allows the session. • Deny—session dropped after the application is identified and there is a rule to block or no rule that allows the session.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Source Address	Original session source IP address.
Destination Address	Original session destination IP address.
NAT Source IP	If Source NAT performed, the post-NAT Source IP address.
NAT Destination IP	If Destination NAT performed, the post-NAT Destination IP address.
Rule Name (Rule)	Name of the rule that the session matched.
Source User	Username of the user who initiated the session.
Destination User	Username of the user to which the session was destined.
Application	Application associated with the session.
Virtual System	Virtual System associated with the session.
Source Zone	Zone the session was sourced from.
Destination Zone	Zone the session was destined to.
Inbound Interface	Interface that the session was sourced from.
Outbound Interface	Interface that the session was destined to.
Log Action	Log Forwarding Profile that was applied to the session.
Session ID	An internal numerical identifier applied to each session.
Repeat Count	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only.
Source Port	Source port utilized by the session.
Destination Port	Destination port utilized by the session.
NAT Source Port	Post-NAT source port.
NAT Destination Port	Post-NAT destination port.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Flags	32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value: • 0x80000000 —session has a packet capture (PCAP) • 0x02000000 —IPv6 session • 0x01000000 —SSL session was decrypted (SSL Proxy) • 0x00800000 —session was denied via URL filtering • 0x00400000 —session has a NAT translation performed (NAT) • 0x00200000 —user information for the session was captured through Captive Portal • 0x00080000 —X-Forwarded-For value from a proxy is in the source user field • 0x00040000 —log corresponds to a transaction within a http proxy session (Proxy Transaction) • 0x00008000 —session is a container page access (Container Page) • 0x00002000 —session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above. • 0x00000800 —symmetric return was used to forward traffic for this session
IP Protocol	IP protocol associated with the session.
Action	Action taken for the session; possible values are: • allow—session was allowed by policy • deny—session was denied by policy • drop—session was dropped silently • drop ICMP—session was silently dropped with an ICMP unreachable message to the host or application • reset both—session was terminated and a TCP reset is sent to both the sides of the connection • reset client—session was terminated and a TCP reset is sent to the client • reset server—session was terminated and a TCP reset is sent to the server
Bytes	Number of total bytes (transmit and receive) for the session.
Bytes Sent	Number of bytes in the client-to-server direction of the session. Available on all models except the PA-4000 Series.
Bytes Received	Number of bytes in the server-to-client direction of the session. Available on all models except the PA-4000 Series.
Packets	Number of total packets (transmit and receive) for the session.
Start Time	Time of session start.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Elapsed Time (sec)	Elapsed time of the session.
Category	URL category associated with the session (if applicable).
Sequence Number	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags	A bit field indicating if the log was forwarded to Panorama.
Source Country	Source country or Internal region for private addresses; maximum length is 32 bytes.
Destination Country	Destination country or Internal region for private addresses. Maximum length is 32 bytes.
Packets Sent (pkts_sent)	Number of client-to-server packets for the session. Available on all models except the PA-4000 Series.
Packets Received (pkts_received)	Number of server-to-client packets for the session. Available on all models except the PA-4000 Series.
Session End Reason (session_end_reason)	The reason a session terminated. If the termination had multiple causes, this field displays only the highest priority reason. The possible session end reason values are as follows, in order of priority (where the first is highest): • threat—The firewall detected a threat associated with a reset, drop, or block (IP address) action. • policy-deny—The session matched a security rule with a deny or drop action. • decrypt-cert-validation—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses client authentication or when the session uses a server certificate with any of the following conditions: expired, untrusted issuer, unknown status, or status verification time-out. This session end reason also displays when the server certificate produces a fatal error alert of type bad_certificate, unsupported_certificate, certificate_revoked, access_denied, or no_certificate_RESERVED (SSLv3 only). • decrypt-unsupport-param—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses an unsupported protocol version, cipher, or SSH algorithm. This session end reason is displays when the session produces a fatal error alert of type unsupported_extension, unexpected_message, or handshake_failure. • decrypt-error—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when firewall resources or the hardware security module (HSM) were unavailable. This session end reason is also

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	displayed when you configured the firewall to block SSL traffic that has SSH errors or that produced any fatal error alert other than those listed for the decrypt-cert-validation and decrypt-unsupport-param end reasons. • tcp-rst-from-client—The client sent a TCP reset to the server. • tcp-rst-from-server—The server sent a TCP reset to the client. • resources-unavailable—The session dropped because of a system resource limitation. For example, the session could have exceeded the number of out-of-order packets allowed per flow or the global out-of-order packet queue. • tcp-fin—One host or both hosts in the connection sent a TCP FIN message to close the session.

	• tcp-reuse—A session is reused and the firewall closes the previous session.
	• decoder—The decoder detects a new connection within the protocol (such as HTTP-Proxy) and ends the previous connection.
	• aged-out—The session aged out.
	• unknown—This value applies in the following situations:
	  • Session terminations that the preceding reasons do not cover (for example, a `clear session all` command).
	  • For logs generated in a PAN-OS release that does not support the session end reason field (releases older than PAN-OS 6.1), the value will be **unknown** after an upgrade to the current PAN-OS release or after the logs are loaded onto the firewall.
	  • In Panorama, logs received from firewalls for which the PAN-OS version does not support session end reasons will have a value of **unknown**.
	• n/a—This value applies when the traffic log type is not **end**.

Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **show readonly dg-meta-data** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Device Name	The hostname of the firewall on which the session was logged.
Action Source (action_source)	Specifies whether the action taken to allow or block an application was defined in the application or in policy. The actions can be allow, deny, drop, reset- server, reset-client or reset-both for the session.
Source VM UUID	Identifies the source universal unique identifier for a guest virtual machine in the VMware NSX environment.
Destination VM UUID	Identifies the destination universal unique identifier for a guest virtual machine in the VMware NSX environment.
Tunnel ID/IMSI	International Mobile Subscriber Identity (IMSI) is a unique number allocated to each mobile subscriber in the GSM/UMTS/EPS system. IMSI shall consist of decimal digits (0 through 9) only and maximum number of digits allowed are 15.
Monitor Tag/IMEI	International Mobile Equipment Identity (IMEI) is a unique 15 or 16 digit number allocated to each mobile station equipment.
Parent Session ID	ID of the session in which this session is tunneled. Applies to inner tunnel (if two levels of tunneling) or inside content (if one level of tunneling) only.
Parent Start Time (parent_start_time)	Year/month/day hours:minutes:seconds that the parent tunnel session began.
Tunnel Type (Tunnel)	Type of tunnel, such as GRE or IPSec.

Threat Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, URL/Filename, Threat ID, Category, Severity, Direction, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Content Type, PCAP_ID, File Digest, Cloud, URL Index, User Agent, File Type, X-Forwarded-For, Referer, Sender, Subject, Recipient, Report ID, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, FUTURE_USE, Source VM UUID, Destination VM UUID, HTTP Method, Tunnel ID/IMSI, Monitor Tag/IMEI, Parent Session ID, Parent Start Time, Tunnel Type, Threat Category, Content Version, FUTURE_USE

Field Name	Description
Receive Time	Time the log was received at the management plane.
Serial Number (serial #)	Serial number of the firewall that generated the log.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Type	Specifies type of log; values are traffic, threat, config, system and hip-match.
Threat/Content Type	Subtype of threat log. Values include the following: • data—Data pattern matching a Data Filtering profile. • file—File type matching a File Blocking profile. • flood—Flood detected via a Zone Protection profile. • packet—Packet-based attack protection triggered by a Zone Protection profile. • scan—Scan detected via a Zone Protection profile. • spyware —Spyware detected via an Anti-Spyware profile. • url—URL filtering log. • virus—Virus detected via an Antivirus profile. • vulnerability —Vulnerability exploit detected via a Vulnerability Protection profile. • wildfire —A WildFire verdict generated when the firewall submits a file to WildFire per a WildFire Analysis profile and a verdict (malicious, phishing, grayware, or benign, depending on what you are logging) is logged in the WildFire Submissions log. • wildfire-virus—Virus detected via an Antivirus profile.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Source Address	Original session source IP address.
Destination Address	Original session destination IP address.
NAT Source IP	If source NAT performed, the post-NAT source IP address.
NAT Destination IP	If destination NAT performed, the post-NAT destination IP address.
Rule Name (rule)	Name of the rule that the session matched.
Source User	Username of the user who initiated the session.
Destination User	Username of the user to which the session was destined.
Application	Application associated with the session.
Virtual System	Virtual System associated with the session.
Source Zone	Zone the session was sourced from.
Destination Zone	Zone the session was destined to.
Inbound Interface	Interface that the session was sourced from.
Outbound Interface	Interface that the session was destined to.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Log Action	Log Forwarding Profile that was applied to the session.
Session ID	An internal numerical identifier applied to each session.
Repeat Count	Number of sessions with same Source IP, Destination IP, Application, and Content/Threat Type seen within 5 seconds; used for ICMP only.
Source Port	Source port utilized by the session.
Destination Port	Destination port utilized by the session.
NAT Source Port	Post-NAT source port.
NAT Destination Port	Post-NAT destination port.
Flags	32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value: • 0x80000000 —session has a packet capture (PCAP) • 0x02000000 —IPv6 session • 0x01000000 —SSL session was decrypted (SSL Proxy) • 0x00800000 —session was denied via URL filtering • 0x00400000 —session has a NAT translation performed (NAT) • 0x00200000 —user information for the session was captured through Captive Portal • 0x00080000 —X-Forwarded-For value from a proxy is in the source user field • 0x00040000 —log corresponds to a transaction within a http proxy session (Proxy Transaction) • 0x00008000 —session is a container page access (Container Page) • 0x00002000 —session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above • 0x00000800 —symmetric return was used to forward traffic for this session
IP Protocol	IP protocol associated with the session.
Action	Action taken for the session; values are alert, allow, deny, drop, drop-all-packets, reset-client, reset-server, reset-both, block-url. • alert—threat or URL detected but not blocked • allow— flood detection alert • deny—flood detection mechanism activated and deny traffic based on configuration • drop— threat detected and associated session was dropped • reset-client —threat detected and a TCP RST is sent to the client • reset-server —threat detected and a TCP RST is sent to the server • reset-both —threat detected and a TCP RST is sent to both the client and the server • block-url —URL request was blocked because it matched a URL category that was set to be blocked

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	• block-ip—threat detected and client IP is blocked • random-drop—flood detected and packet was randomly dropped • sinkhole—DNS sinkhole activated • syncookie-sent—syncookie alert • block-continue (URL subtype only)—a HTTP request is blocked and redirected to a Continue page with a button for confirmation to proceed • continue (URL subtype only)—response to a block-continue URL continue page indicating a block-continue request was allowed to proceed • block-override (URL subtype only)—a HTTP request is blocked and redirected to an Admin override page that requires a pass code from the firewall administrator to continue • override-lockout (URL subtype only)—too many failed admin override pass code attempts from the source IP. IP is now blocked from the block-override redirect page • override (URL subtype only)—response to a block-override page where a correct pass code is provided and the request is allowed • block (Wildfire only)—file was blocked by the firewall and uploaded to Wildfire
URL/Filename	Field with variable length with a maximum of 1023 characters The actual URI when the subtype is URL File name or file type when the subtype is file File name when the subtype is virus File name when the subtype is WildFire
Threat Content Name	Palo Alto Networks identifier for the threat. It is a description string followed by a 64-bit numerical identifier in parentheses for some Subtypes: • 8000 – 8099— scan detection • 8500 – 8599— flood detection • 9999— URL filtering log • 10000 – 19999 —spyware phone home detection • 20000 – 29999 —spyware download detection • 30000 – 44999 —vulnerability exploit detection • 52000 – 52999— filetype detection • 60000 – 69999 —data filtering detection  *Threat ID ranges for virus detection, WildFire signature feed, and DNS C2 signatures used in previous releases have been replaced with permanent, globally unique IDs. Refer to the Threat/Content Type and Threat Category (thr_category) field names to create updated reports, filter threat logs, and ACC activity.*

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Category	For URL Subtype, it is the URL Category; For WildFire subtype, it is the verdict on the file and is either 'malicious', 'phishing', 'grayware', or 'benign'; For other subtypes, the value is 'any'.
Severity	Severity associated with the threat; values are informational, low, medium, high, critical.
Direction	Indicates the direction of the attack, client-to-server or server-to-client: • 0—direction of the threat is client to server • 1—direction of the threat is server to client
Sequence Number	A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space.
Action Flags	A bit field indicating if the log was forwarded to Panorama.
Source Country	Source country or Internal region for private addresses. Maximum length is 32 bytes.
Destination Country	Destination country or Internal region for private addresses. Maximum length is 32 bytes.
Content Type (contenttype)	Applicable only when Subtype is URL. Content type of the HTTP response data. Maximum length 32 bytes.
PCAP ID (pcap_id)	The packet capture (pcap) ID is a 64 bit unsigned integral denoting an ID to correlate threat pcap files with extended pcaps taken as a part of that flow. All threat logs will contain either a pcap_id of 0 (no associated pcap), or an ID referencing the extended pcap file.
File Digest (filedigest)	Only for WildFire subtype; all other types do not use this field The filedigest string shows the binary hash of the file sent to be analyzed by the WildFire service.
Cloud (cloud)	Only for WildFire subtype; all other types do not use this field. The cloud string displays the FQDN of either the WildFire appliance (private) or the WildFire cloud (public) from where the file was uploaded for analysis.
URL Index (url_idx)	Used in URL Filtering and WildFire subtypes. When an application uses TCP keepalives to keep a connection open for a length of time, all the log entries for that session have a single session ID. In such cases, when you have a single threat log (and session ID) that includes multiple URL entries, the url_idx is a counter that allows you to correlate the order of each log entry within the single session. For example, to learn the URL of a file that the firewall forwarded to WildFire for analysis, locate the session ID and the url_idx from the WildFire Submissions log and search for the same session ID and url_idx

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	in your URL filtering logs. The log entry that matches the session ID and url_idx will contain the URL of the file that was forwarded to WildFire.
User Agent (user_agent)	Only for the URL Filtering subtype; all other types do not use this field. The User Agent field specifies the web browser that the user used to access the URL, for example Internet Explorer. This information is sent in the HTTP request to the server.
File Type (filetype)	Only for WildFire subtype; all other types do not use this field. Specifies the type of file that the firewall forwarded for WildFire analysis.
X-Forwarded-For (xff)	Only for the URL Filtering subtype; all other types do not use this field. The X-Forwarded-For field in the HTTP header contains the IP address of the user who requested the web page. It allows you to identify the IP address of the user, which is useful particularly if you have a proxy server on your network that replaces the user IP address with its own address in the source IP address field of the packet header.
Referer (referer)	Only for the URL Filtering subtype; all other types do not use this field. The Referer field in the HTTP header contains the URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested.
Sender (sender)	Only for WildFire subtype; all other types do not use this field. Specifies the name of the sender of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall.
Subject (subject)	Only for WildFire subtype; all other types do not use this field. Specifies the subject of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall.
Recipient (recipient)	Only for WildFire subtype; all other types do not use this field. Specifies the name of the receiver of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall.
Report ID (reportid)	Only for WildFire subtype; all other types do not use this field. Identifies the analysis request on the WildFire cloud or the WildFire appliance.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods:

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Source VM UUID	Identifies the source universal unique identifier for a guest virtual machine in the VMware NSX environment.
Destination VM UUID	Identifies the destination universal unique identifier for a guest virtual machine in the VMware NSX environment.
HTTP Method	Only in URL filtering logs. Describes the HTTP Method used in the web request. Only the following methods are logged: Connect, Delete, Get, Head, Options, Post, Put.
Tunnel ID/IMSI	International Mobile Subscriber Identity (IMSI) is a unique number allocated to each mobile subscriber in the GSM/UMTS/EPS system. IMSI shall consist of decimal digits (0 through 9) only and maximum number of digits allowed are 15.
Monitor Tag/IMEI	International Mobile Equipment Identity (IMEI) is a unique 15 or 16 digit number allocated to each mobile station equipment.
Parent Session ID	ID of the session in which this session is tunneled. Applies to inner tunnel (if two levels of tunneling) or inside content (if one level of tunneling) only.
Parent Start Time (parent_start_time)	Year/month/day hours:minutes:seconds that the parent tunnel session began.
Tunnel Type (Tunnel)	Type of tunnel, such as GRE or IPSec.
Threat Category (thr_category)	Describes threat categories used to classify different types of threat signatures.
Content Version (contentver)	Applications and Threats version on your firewall when the log was generated.

HIP Match Log Fields

If you are running PAN-OS 8.0.5 through PAN-OS 8.0.10, refer to the HIP Match Log Fields in PAN-OS 8.0.5 Through 8.0.10 on page 397.

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source User, Virtual System, Machine name, OS, Source Address, HIP, Repeat Count, HIP Type,

© 2018 Palo Alto Networks, Inc.

FUTURE_USE, FUTURE_USE, Sequence Number, Action Flags, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, IPv6 Source Address, Host ID

Field Name	Description
Receive Time	Time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the firewall that generated the log.
Type	Type of log; values are traffic, threat, config, system and hip-match.
Threat/Content Type	Subtype of HIP match log; unused.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Source User	Username of the user who initiated the session.
Virtual System	Virtual System associated with the HIP match log.
Machine Name (machinename)	Name of the user's machine.
OS	The operating system installed on the user's machine or device (or on the client system).
Source Address	IP address of the source user.
HIP (matchname)	Name of the HIP object or profile.
Repeat Count	Number of times the HIP profile matched.
HIP Type (matchtype)	Whether the hip field represents a HIP object or a HIP profile.
Sequence Number	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags	A bit field indicating if the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **`show readonly dg-meta-data`**

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name	The hostname of the firewall on which the session was logged.
Virtual System ID	A unique identifier for a virtual system on a Palo Alto Networks firewall.
IPv6 System Address	IPv6 address of the user's machine or device.
Host ID (hostid) —8.0.11 only	Unique ID GlobalProtect assigns to identify the host.

HIP Match Log Fields in PAN-OS 8.0.5 Through 8.0.10

Format: FUTURE_USE, Receive Time, Serial Number, Sequence Number, Action Flags, Type, Threat/Content Type, FUTURE_USE, Generated Time, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, Source User, Virtual System, Machine Name, Source Address, HIP, Repeat Count, OS, HIP Type, FUTURE_USE, FUTURE_USE, IPv6 Source Address, Host ID

Field Name	Description
Receive Time	Time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the firewall that generated the log.
Sequence Number	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags	A bit field indicating if the log was forwarded to Panorama.
Type	Type of log; values are traffic, threat, config, system and hip-match.
Threat/Content Type	Subtype of HIP match log; unused.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name	The hostname of the firewall on which the session was logged.
Virtual System ID	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Source User	Username of the user who initiated the session.
Virtual System	Virtual System associated with the HIP match log.
Machine Name (machinename)	Name of the user's machine.
Source Address	IP address of the source user.
HIP (matchname)	Name of the HIP object or profile.
Repeat Count	Number of times the HIP profile matched.
OS	The operating system installed on the user's machine or device (or on the client system).
HIP Type (matchtype)	Whether the hip field represents a HIP object or a HIP profile.
IPv6 System Address	IPv6 address of the user's machine or device.
Host ID (hostid)	Unique ID GlobalProtect assigns to identify the host.

User-ID Log Fields

If you are running PAN-OS 8.0 through PAN-OS 8.0.10, refer to the User-ID Log Fields Pre 8.0.11 on page 400.

Format: FUTURE_USER, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Virtual System, Source IP, User, Data Source Name, Event ID, Repeat Count, Time Out Threshold, Source Port, Destination Port, Data Source, Data Source Type, Sequence Number, Action Flags, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, Factor Type, Factor Completion Time, Factor Number, User Group Flags, User Provided by Source

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Receive Time (receive_time)	Time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the firewall that generated the log.
Type (type)	Specifies type of log; values are traffic, threat, config, system and hip-match.
Threat/Content Type	Subtype of traffic log; values are start, end, drop, and deny "Start-session started "End-session ended "Drop-session dropped before the application is identified and there is no rule that allows the session. "Deny-session dropped after the application is identified and there is a rule to block or no rule that allows the session.
Generated Time (Generate Time)	The time the log was generated on the dataplane.
Virtual System	Virtual System associated with the configuration log.
Source IP	Original session source IP address.
User	Identifies the end user.
Data Source Name	User-ID source that sends the IP (Port)-User Mapping.
Event ID	String showing the name of the event.
Repeat Count	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only.
Time Out Threshold (timeout)	Timeout after which the IP/User Mappings are cleared.
Source Port (beginport)	Source port utilized by the session.
Destination Port (endport)	Destination port utilized by the session.
Data Source	Source from which mapping information is collected.
Data Source Type	Mechanism used to identify the IP/User mappings within a data source.
Sequence Number	Serial number of the firewall that generated the log.
Action Flags	A bit field indicating if the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **show readonly dg-meta-data** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name	The hostname of the firewall on which the session was logged.
Virtual System ID	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Factor Type	Vendor used to authenticate a user when Multi Factor authentication is present.
Factor Completion Time	Time the authentication was completed.
Factor Number	Indicates the use of primary authentication (1) or additional factors (2, 3).
User Group Flags (ugflags)	Displays whether the user group was found during user group mapping. Supported values are: • User Group Found—Indicates whether the user could be mapped to a group. • Duplicate User—Indicates whether duplicate users were found in a user group. Displays N/A if no user group is found.
User Provided by Source (userbysource)	Indicates the username received from source through IP-User mapping.

User-ID Log Fields Pre 8.0.11

Format: FUTURE_USER, Receive Time, Serial Number, Sequence Number, Action Flags, Type, Threat/Content Type, FUTURE_USE, Generated Time, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, Virtual System, Source IP, User, Data Source Name, Event ID, Repeat Count, Time Out Threshold, Source Port, Destination Port, Data Source, Data Source Type, FUTURE_USE, FUTURE_USE, Factor Type, Factor Completion Time, Factor Number

Field Name	Description
Receive Time (receive_time)	Time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the firewall that generated the log.
Sequence Number	Serial number of the firewall that generated the log.
Action Flags	A bit field indicating if the log was forwarded to Panorama.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Type (type)	Specifies type of log; values are traffic, threat, config, system and hip-match.
Threat/Content Type	Subtype of traffic log; values are start, end, drop, and deny "Start-session started "End-session ended "Drop-session dropped before the application is identified and there is no rule that allows the session. "Deny-session dropped after the application is identified and there is a rule to block or no rule that allows the session.
Generated Time (Generate Time)	The time the log was generated on the dataplane.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **show readonly dg-meta-data** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name	The hostname of the firewall on which the session was logged.
Virtual System ID	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Virtual System	Virtual System associated with the configuration log.
Source IP	Original session source IP address.
User	Identifies the end user.
Data Source Name	User-ID source that sends the IP (Port)-User Mapping.
Event ID	String showing the name of the event.
Repeat Count	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only.
Time Out Threshold (timeout)	Timeout after which the IP/User Mappings are cleared.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Source Port (beginport)	Source port utilized by the session.
Destination Port (endport)	Destination port utilized by the session.
Data Source	Source from which mapping information is collected.
Data Source Type	Mechanism used to identify the IP/User mappings within a data source.
Factor Type	Vendor used to authenticate a user when Multi Factor authentication is present.
Factor Completion Time	Time the authentication was completed.
Factor Number	Indicates the use of primary authentication (1) or additional factors (2, 3).

Tunnel Inspection Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Severity, Sequence Number, Action Flags, Source Location, Destination Location, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Tunnel ID/IMSI, Monitor Tag/IMEI, Parent Session ID, Parent Start Time, Tunnel, Bytes, Bytes Sent, Bytes Received, Packets, Packets Sent, Maximum Encapsulation, Unknown Protocol, Strict Check, Tunnel Fragment, Sessions Created, Sessions Closed, Session End Reason, Action Source, Start Time, Elapsed Time

Field Name	Description
Receive Time	Month, day, and time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the firewall that generated the log.
Type	Type of log as it pertains to the session: start or end.
Threat/Content Type	Subtype of traffic log; values are start, end, drop, and deny • Start—session started • End—session ended • Drop—session dropped before the application is identified and there is no rule that allows the session. • Deny—session dropped after the application is identified and there is a rule to block or no rule that allows the session.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Source Address	Source IP address of packets in the session.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Destination Address	Destination IP address of packets in the session.
NAT Source IP	If Source NAT performed, the post-NAT Source IP address.
NAT Destination IP	If Destination NAT performed, the post-NAT Destination IP address.
Rule Name (Rule)	Name of the Security policy rule in effect on the session.
Source User	Source User ID of packets in the session.
Destination User	Destination User ID of packets in the session.
Application	Tunneling protocol used in the session.
Virtual System	Virtual System associated with the session.
Source Zone	Source zone of packets in the session.
Destination Zone	Destination zone of packets in the session.
Inbound Interface	Interface that the session was sourced from.
Outbound Interface	Interface that the session was destined to.
Log Action	Log Forwarding Profile that was applied to the session.
Session ID	Session ID of the session being logged.
Repeat Count	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only.
Source Port	Source port utilized by the session.
Destination Port	Destination port utilized by the session.
NAT Source Port	Post-NAT source port.
NAT Destination Port	Post-NAT destination port.
Flags	32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value: • 0x80000000 —session has a packet capture (PCAP) • 0x02000000 —IPv6 session • 0x01000000 —SSL session was decrypted (SSL Proxy) • 0x00800000 —session was denied via URL filtering • 0x00400000 —session has a NAT translation performed (NAT) • 0x00200000 —user information for the session was captured via the captive portal (Captive Portal) • 0x00080000 —X-Forwarded-For value from a proxy is in the source user field

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	• 0x00040000 —log corresponds to a transaction within a http proxy session (Proxy Transaction) • 0x00008000 —session is a container page access (Container Page) • 0x00002000 —session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above. • 0x00000800 —symmetric return was used to forward traffic for this session
Protocol (IP Protocol)	IP protocol associated with the session.
Action	Action taken for the session; possible values are: • Allow—session was allowed by policy • Deny—session was denied by policy • Drop—session was dropped silently • Drop ICMP—session was silently dropped with an ICMP unreachable message to the host or application • Reset both—session was terminated and a TCP reset is sent to both the sides of the connection • Reset client—session was terminated and a TCP reset is sent to the client • Reset server—session was terminated and a TCP reset is sent to the server
Severity	Severity associated with the event; values are informational, low, medium, high, critical.
Sequence Number	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7000 Series firewalls.
Action Flags	A bit field indicating if the log was forwarded to Panorama.
Source Location (source country)	Source country or Internal region for private addresses; maximum length is 32 bytes.
Destination Location (destination country)	Destination country or Internal region for private addresses. Maximum length is 32 bytes.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name	The hostname of the firewall on which the session was logged.
Tunnel ID/IMSI	ID of the tunnel being inspected or the International Mobile Subscriber Identity (IMSI) ID of the mobile user.
Monitor Tag/IMEI	Monitor name you configured for the Tunnel Inspection policy rule or the International Mobile Equipment Identity (IMEI) ID of the mobile device.
Parent Session ID	ID of the session in which this session is tunneled. Applies to inner tunnel (if two levels of tunneling) or inside content (if one level of tunneling) only.
Parent Start Time (parent_start_time)	Year/month/day hours:minutes:seconds that the parent tunnel session began.
Tunnel Type (Tunnel)	Type of tunnel, such as GRE or IPSec.
Bytes	Number of bytes in the session.
Bytes Sent	Number of bytes in the client-to-server direction of the session.
Bytes Received	Number of bytes in the server-to-client direction of the session.
Packets	Number of total packets (transmit and receive) for the session.
Packets Sent (pkts_sent)	Number of client-to-server packets for the session. Available on all models except the PA-4000 Series.
Packets Received (pkts_received)	Number of server-to-client packets for the session. Available on all models except the PA-4000 Series.
Maximum Encapsulation (max_encap)	Number of packets the firewall dropped because the packet exceeded the maximum number of encapsulation levels configured in the Tunnel Inspection policy rule (Drop packet if over maximum tunnel inspection level).
Unknown Protocol (unknown_proto)	Number of packets the firewall dropped because the packet contains an unknown protocol, as enabled in the Tunnel Inspection policy rule (Drop packet if unknown protocol inside tunnel).
Strict Checking (strict_check)	Number of packets the firewall dropped because the tunnel protocol header in the packet failed to comply with the RFC for the tunnel protocol, as enabled in the Tunnel Inspection policy rule (**Drop packet if tunnel protocol fails strict header check**).
Tunnel Fragment (tunnel_fragment)	Number of packets the firewall dropped because of fragmentation errors.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Sessions Created (sessions_created)	Number of inner sessions created.
Sessions Closed (sessions_closed)	Number of completed/closed sessions created.
Session End Reason (session_end_reason)	The reason a session terminated. If the termination had multiple causes, this field displays only the highest priority reason. The possible session end reason values are as follows, in order of priority (where the first is highest): • threat—The firewall detected a threat associated with a reset, drop, or block (IP address) action. • policy-deny—The session matched a security rule with a deny or drop action. • decrypt-cert-validation—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses client authentication or when the session uses a server certificate with any of the following conditions: expired, untrusted issuer, unknown status, or status verification time-out. This session end reason also displays when the server certificate produces a fatal error alert of type bad_certificate, unsupported_certificate, certificate_revoked, access_denied, or no_certificate_RESERVED (SSLv3 only). • decrypt-unsupport-param—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses an unsupported protocol version, cipher, or SSH algorithm. This session end reason is displays when the session produces a fatal error alert of type unsupported_extension, unexpected_message, or handshake_failure. • decrypt-error—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when firewall resources or the hardware security module (HSM) were unavailable. This session end reason is also displayed when you configured the firewall to block SSL traffic that has SSH errors or that produced any fatal error alert other than those listed for the decrypt-cert-validation and decrypt-unsupport-param end reasons. • tcp-rst-from-client—The client sent a TCP reset to the server. • tcp-rst-from-server—The server sent a TCP reset to the client. • resources-unavailable—The session dropped because of a system resource limitation. For example, the session could have exceeded the number of out-of-order packets allowed per flow or the global out-of-order packet queue. • tcp-fin—One host or both hosts in the connection sent a TCP FIN message to close the session. • tcp-reuse—A session is reused and the firewall closes the previous session. • decoder—The decoder detects a new connection within the protocol (such as HTTP-Proxy) and ends the previous connection. • aged-out—The session aged out. • unknown—This value applies in the following situations: • Session terminations that the preceding reasons do not cover (for example, a `clear session all` command).

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	• For logs generated in a PAN-OS release that does not support the session end reason field (releases older than PAN-OS 6.1), the value will be `unknown` after an upgrade to the current PAN-OS release or after the logs are loaded onto the firewall. • In Panorama, logs received from firewalls for which the PAN-OS version does not support session end reasons will have a value of `unknown`. • n/a—This value applies when the traffic log type is not **`end`**.
Action Source (action_source)	Specifies whether the action taken to allow or block an application was defined in the application or in policy. The actions can be allow, deny, drop, reset-server, reset-client or reset-both for the session.
Start Time (start)	Year/month/day hours:minutes:seconds that the session began.
Elapsed Time (sec)	Elapsed time of the session.

Authentication Log Fields

If you are running PAN-OS 8.0.5 through PAN-OS 8.0.10, refer to the Authentication Log Fields in PAN-OS 8.0.5 Through 8.0.10 on page 409.

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Virtual System, Source IP, User, Normalize User, Object, Authentication Policy, Repeat Count, Authentication ID, Vendor, Log Action, Server Profile, Description, Client Type, Event Type, Factor Number, Sequence Number, Action Flags, Device Group Hierarchy 1, Device Group Hierarchy 2, Device Group Hierarchy 3, Device Group Hierarchy 4, Virtual System Name, Device Name, Virtual System ID, Authentication Protocol

Field Name	Description
Receive Time	Time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the device that generated the log.
Type	Type of log; values are traffic, threat, config, system and hip-match.
Threat/Content Type	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Virtual System	Virtual System associated with the session.
Source IP	Original session source IP address.
User	End user being authenticated.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Normalize User	Normalized version of username being authenticated (such as appending a domain name to the username).
Object	Name of the object associated with the system event.
Authentication Policy	Policy invoked for authentication before allowing access to a protected resource.
Repeat Count	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only.
Authentication ID	Unique ID given across primary authentication and additional (multi factor) authentication.
Vendor	Vendor providing additional factor authentication.
Log Action	Log Forwarding Profile that was applied to the session.
Server Profile (serverprofile)	Authentication server used for authentication.
Description (desc)	Additional authentication information.
Client Type	Type of client used to complete authentication (such as authentication portal).
Event Type	Result of the authentication attempt.
Factor Number	Indicates the use of primary authentication (1) or additional factors (2, 3).
Sequence Number	A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space.
Action Flags	A bit field indicating if the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Device Name	The hostname of the firewall on which the session was logged.
Virtual System ID —8.0.11 only	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Authentication Protocol (authproto) —8.0.11 only	Indicates the authentication protocol used by the server. For example, PEAP with GTC.

Authentication Log Fields in PAN-OS 8.0.5 Through 8.0.10

Format: FUTURE_USE, Receive Time, Serial Number, Sequence Number, Action Flags, Type, Threat/Content Type, FUTURE_USE, Generated Time, Device Group Hierarchy 1, Device Group Hierarchy 2, Device Group Hierarchy 3, Device Group Hierarchy 4, Virtual System Name, Device Name, Virtual System ID, Virtual System, Source IP, User, Normalize User, Object, Authentication Policy, Authentication ID, Vendor, Log Action, Repeat Count, Server Profile, Description, Client Type, Event Type, Factor Number, Authentication Protocol

Field Name	Description
Receive Time	Time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the device that generated the log.
Sequence Number	A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space.
Action Flags	A bit field indicating if the log was forwarded to Panorama.
Type	Type of log; values are traffic, threat, config, system and hip-match.
Threat/Content Type	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **`show readonly dg-meta-data`**

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name	The hostname of the firewall on which the session was logged.
Virtual System ID	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Virtual System	Virtual System associated with the session.
Source IP	Original session source IP address.
User	End user being authenticated.
Normalize User	Normalized version of username being authenticated (such as appending a domain name to the username).
Object	Name of the object associated with the system event.
Authentication Policy	Policy invoked for authentication before allowing access to a protected resource.
Authentication ID	Unique ID given across primary authentication and additional (multi factor) authentication.
Vendor	Vendor providing additional factor authentication.
Log Action	Log Forwarding Profile that was applied to the session.
Repeat Count	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only.
Server Profile (serverprofile)	Authentication server used for authentication.
Description (desc)	Additional authentication information.
Client Type	Type of client used to complete authentication (such as authentication portal).
Event Type	Result of the authentication attempt.
Factor Number	Indicates the use of primary authentication (1) or additional factors (2, 3).
Authentication Protocol (authproto)	Indicates the authentication protocol used by the server. For example, PEAP with GTC.

© 2018 Palo Alto Networks, Inc.

Config Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Host, Virtual System, Command, Admin, Client, Result, Configuration Path, Before Change Detail, After Change Detail, Sequence Number, Action Flags, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name

Field Name	Description
Receive Time	Time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the device that generated the log.
Type	Type of log; values are traffic, threat, config, system and hip-match.
Subtype	Subtype of the configuration log; unused.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Host	Hostname or IP address of the client machine
Virtual System	Virtual System associated with the configuration log
Command (cmd)	Command performed by the Admin; values are add, clone, commit, delete, edit, move, rename, set.
Admin (admin)	Username of the Administrator performing the configuration
Client (client)	Client used by the Administrator; values are Web and CLI
Result (result)	Result of the configuration action; values are Submitted, Succeeded, Failed, and Unauthorized
Configuration Path (path)	The path of the configuration command issued; up to 512 bytes in length
Before Change Detail (before_change_detail)	This field is in custom logs only; it is not in the default format. It contains the full xpath before the configuration change.
After Change Detail (after_change_detail)	This field is in custom logs only; it is not in the default format. It contains the full xpath after the configuration change.
Sequence Number (seqno)	A 64bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Device Group Hierarchy	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating

© 2018 Palo Alto Networks, Inc.

Field Name	Description
(dg_hier_level_1 to dg_hier_level_4)	the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name	The hostname of the firewall on which the session was logged.

System Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Content/Threat Type, FUTURE_USE, Generated Time, Virtual System, Event ID, Object, FUTURE_USE, FUTURE_USE, Module, Severity, Description, Sequence Number, Action Flags, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name

Field Name	Description
Receive Time	Time the log was received at the management plane
Serial Number (Serial #)	Serial number of the firewall that generated the log
Type	Type of log; values are traffic, threat, config, system and hip-match
Content/Threat Type	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
Generated Time (Generate Time)	Time the log was generated on the dataplane
Virtual System	Virtual System associated with the configuration log
Event ID	String showing the name of the event
Object	Name of the object associated with the system event

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Module (module)	This field is valid only when the value of the Subtype field is general. It provides additional information about the sub-system generating the log; values are general, management, auth, ha, upgrade, chassis
Severity	Severity associated with the event; values are informational, low, medium, high, critical
Description	Detailed description of the event, up to a maximum of 512 bytes
Sequence Number	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
Action Flags	A bit field indicating if the log was forwarded to Panorama
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name	The hostname of the firewall on which the session was logged.

Correlated Events Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Content/Threat Type, FUTURE_USE, Generated Time, Source Address. Source User, Virtual System, Category, Severity, Device Group Hierarchy Level 1, Device Group Hierarchy Level 2, Device Group Hierarchy Level 3, Device Group Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID, Object Name, Object ID, Evidence

Field Name	Description
Receive Time	Time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the device that generated the log.
Type	Type of log; values are traffic, threat, config, system and hip-match.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Content/Threat Type	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Source Address	IP address of the user who initiated the event.
Source User	Username of the user who initiated the event.
Virtual System	Virtual System associated with the configuration log.
Category	A summary of the kind of threat or harm posed to the network, user, or host.
Severity	Severity associated with the event; values are informational, low, medium, high, critical.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicate the device group's location within a device group hierarchy. The firewall (or virtual system) generating the log includes the identification number of each ancestor in its device group hierarchy. The shared device group (level 0) is not included in this structure. If the log values are 12, 34, 45, 0, it means that the log was generated by a firewall (or virtual system) that belongs to device group 45, and its ancestors are 34, and 12. To view the device group names that correspond to the value 12, 34 or 45, use one of the following methods: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name	The hostname of the firewall on which the session was logged.
Virtual System ID	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Object Name (objectname)	Name of the correlation object that was matched on.
Object ID	Name of the object associated with the system event.
Evidence	A summary statement that indicates how many times the host has matched against the conditions defined in the correlation object. For example, Host visited known malware URl (19 times).

© 2018 Palo Alto Networks, Inc.

GTP Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, FUTURE_USE, FUTURE_USE, FUTURE_USE, Generated Time, Source Address, Destination Address, FUTURE_USE, FUTURE_USE, Rule Name, FUTURE_USE, FUTURE_USE, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, FUTURE_USE, Source Port, Destination Port, FUTURE_USE, FUTURE_USE, FUTURE_USE, Protocol, Action, GTP Event Type, MSISDN, Access Point Name, Radio Access Technology, GTP Message Type, End User IP Address, Tunnel Endpoint Identifier1, Tunnel Endpoint Identifier2, GTP Interface, GTP Cause, Severity, Serving Country MCC, Serving Network MNC, Area Code, Cell ID, GTP Event Code, FUTURE_USE, FUTURE_USE, Source Location, Destination Location, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, Tunnel ID/IMSI, Monitor Tag/IMEI, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE

Field Name	Description
Receive Time	Month, Day and time the log was received at the management plane.
Serial Number (Serial #)	Serial number of the firewall that generated the log.
Generated Time (Generate Time)	Time the log was generated on the dataplane.
Source Address	Source IP address of packets in the session.
Destination Address	Destination IP address of packets in the session.
Rule Name (Rule)	Name of the Security policy rule in effect on the session.
Application	Tunneling protocol used in the session.
Virtual System	Virtual System associated with the session.
Source Zone	Source zone of packets in the session.
Destination Zone	Destination zone of packets in the session.
Inbound Interface	Interface that the session was sourced from.
Outbound Interface	Interface that the session was destined to.
Log Action	Log Forwarding Profile that was applied to the session.
Session ID	Session ID of the session being logged.
Source Port	Source port utilized by the session.
Destination Port	Destination port utilized by the session.
Protocol (IP Protocol)	IP protocol associated with the session.
Action	Action taken for the session; possible values are:

© 2018 Palo Alto Networks, Inc.

Field Name	Description
	• allow—session was allowed by policy • deny—session was denied by policy
GTP Event Type (event_type)	Defines event triggered by a GTP message when checks in GTP protection profile are applied to the GTP traffic. Also triggered by the start or end of a GTP session.
MSISDN (msisdn)	Service identity associated with the mobile subscriber composed of a Country Code, National Destination Code and a Subscriber. Consists of decimal digits (0-9) only with a maximum of 15 digits.
Access Point Name (apn)	Reference to a Packet Data Network Data Gateway (PGW)/ Gateway GPRS Support Node in a mobile network. Composed of a mandatory APN Network Identifier and an optional APN Operator Identifier.
Radio Access Technology (rat)	Type of technology used for radio access. For example, EUTRAN, WLAN, Virtual, HSPA Evolution, GAN and GERAN.
GTP Message Type (msg_type)	Indicates the GTP message type.
End IP Address (end_ip_adr)	IP address of a mobile subscriber allocated by a PGW/GGSN.
Tunnel Endpoint Identifier1 (teid1)	Identifies the GTP tunnel in the network node. TEID1 is the first TEID in the GTP message.
Tunnel Endpoint Identifier2 (teid2)	Identifies the GTP tunnel in the network node. TEID2 is the second TEID in the GTP message.
GTP Interface (gtp_interface)	3GPP interface from which a GTP message is received.
GTP Cause (cause_code)	GTP cause value in logs responses which contain an Information Element that provides information about acceptance or rejection of GTP requests by a network node.
Severity	Severity associated with the event; values are informational, low, medium, high, critical.
Serving Network MCC (mcc)	Mobile country code of serving core network operator.
Serving Network MNC (mnc)	Mobile network code of serving core network operator.
Area Code (area_code)	Area within a Public Land Mobile Network (PLMN).
Cell ID (cell_id)	Base station within an area code.
GTP Event Code (event_code)	Event code describing the GTP event.
Source Location (source country)	Source country or Internal region for private addresses; maximum length is 32 bytes.

© 2018 Palo Alto Networks, Inc.

Field Name	Description
Destination Location (destination country)	Destination country or Internal region for private addresses; maximum length is 32 bytes.
Tunnel ID/IMSI	International Mobile Subscriber Identity (IMSI) is a unique number allocated to each mobile subscriber in the GSM/UMTS/EPS system. IMSI shall consist of decimal digits (0 through 9) only and maximum number of digits allowed are 15.
Monitor Tag/IMEI	International Mobile Equipment Identity (IMEI) is a unique 15 or 16 digit number allocated to each mobile station equipment.

Syslog Severity

The syslog severity is set based on the log type and contents.

Log Type/Severity	Syslog Severity
Traffic	Info
Config	Info
Threat/System—Informational	Info
Threat/System—Low	Notice
Threat/System—Medium	Warning
Threat/System—High	Warning
Threat/System—Critical	Critical

Custom Log/Event Format

To facilitate the integration with external log parsing systems, the firewall allows you to customize the log format; it also allows you to add custom *Key: Value* attribute pairs. Custom message formats can be configured under **Device** > **Server Profiles** > **Syslog** > **Syslog Server Profile** > **Custom Log Format**.

To achieve ArcSight Common Event Format (CEF) compliant log formatting, refer to the CEF Configuration Guide.

Escape Sequences

Any field that contains a comma or a double-quote is enclosed in double quotes. Furthermore, if a double-quote appears inside a field it is escaped by preceding it with another double-quote. To maintain backward compatibility, the Misc field in threat log is always enclosed in double-quotes.

© 2018 Palo Alto Networks, Inc.

SNMP Monitoring and Traps

The following topics describe how Palo Alto Networks firewalls, Panorama, and WF-500 appliances implement SNMP, and the procedures to configure SNMP monitoring and trap delivery.

- SNMP Support on page 418
- Use an SNMP Manager to Explore MIBs and Objects on page 419
- Enable SNMP Services for Firewall-Secured Network Elements on page 421
- Monitor Statistics Using SNMP on page 422
- Forward Traps to an SNMP Manager on page 423
- Supported MIBs on page 425

SNMP Support

You can use an SNMP manager to monitor event-driven alerts and operational statistics for the firewall, Panorama, or WF-500 appliance and for the traffic they process. The statistics and traps can help you identify resource limitations, system changes or failures, and malware attacks. You configure alerts by forwarding log data as traps, and enable the delivery of statistics in response to GET messages (requests) from your SNMP manager. Each trap and statistic has an object identifier (OID). Related OIDs are organized hierarchically within the Management Information Bases (MIBs) that you load into the SNMP manager to enable monitoring.



When an event triggers SNMP trap generation (for example, an interface goes down), the firewall, Panorama virtual appliance, M-Series appliance, and WF-500 appliance respond by updating the corresponding SNMP object (for example, the interfaces MIB) instead of waiting for the periodic update of all objects that occurs every ten seconds. This ensures that your SNMP manager displays the latest information when polling an object to confirm an event.

The firewall, Panorama, and WF-500 appliance support SNMP Version 2c and Version 3. Decide which to use based on the version that other devices in your network support and on your network security requirements. SNMPv3 is more secure and enables more granular access control for system statistics than SNMPv2c. The following table summarizes the security features of each version. You select the version and configure the security features when you Monitor Statistics Using SNMP and Forward Traps to an SNMP Manager.

SNMPVer	Authentication	Message Privacy	Message	MIB Access Granularity
SNMPv2c	Community string	No (cleartext)	No	SNMP community access for all MIBs on a device
SNMPv3	EngineID, username, and authentication password (SHA hashing for the password)	Privacy password for AES 128 encryption of SNMP messages	Yes	User access based on views that include or exclude specific OIDs

SNMP Implementation illustrates a deployment in which firewalls forward traps to an SNMP manager while also forwarding logs to Log Collectors. Alternatively, you could configure the Log Collectors to forward the firewall traps to the SNMP manager. For details on these deployments, refer to Log Forwarding Options in Centralized Logging and Reporting. In all deployments, the SNMP manager gets statistics directly from the

© 2018 Palo Alto Networks, Inc.

firewall, Panorama, or WF-500 appliance. In this example, a single SNMP manager collects both traps and statistics, though you can use separate managers for these functions if that better suits your network.

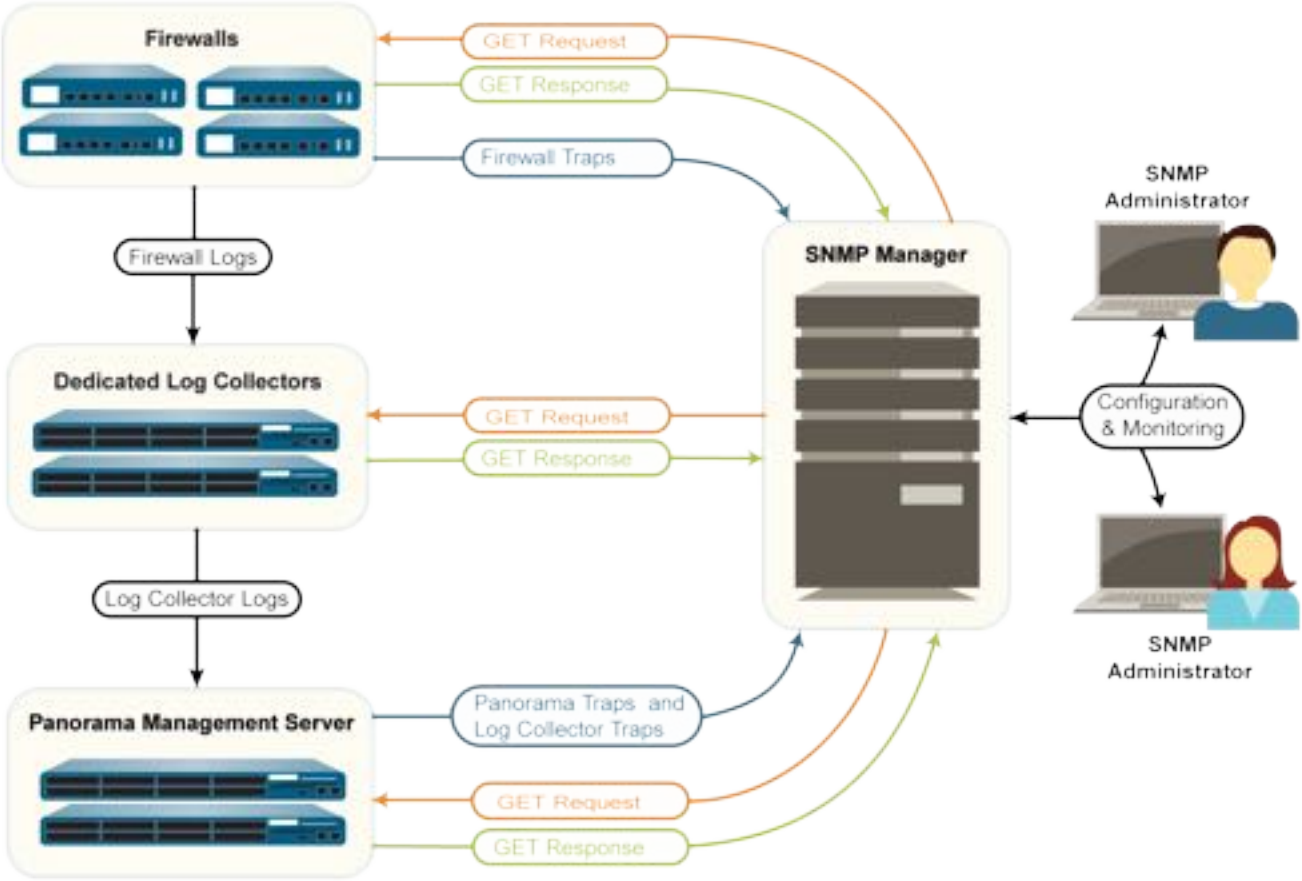


Figure 2: SNMP Implementation

Use an SNMP Manager to Explore MIBs and Objects

To use SNMP for monitoring Palo Alto Networks firewalls, Panorama, or WF-500 appliances, you must first load the Supported MIBs on page 425 into your SNMP manager and determine which object identifiers (OIDs) correspond to the system statistics and traps you want to monitor. The following topics provide an overview of how to find OIDs and MIBs in an SNMP manager. For the specific steps to perform these tasks, refer to your SNMP management software.

- Identify a MIB Containing a Known OID on page 419
- Walk a MIB on page 420
- Identify the OID for a System Statistic or Trap on page 421

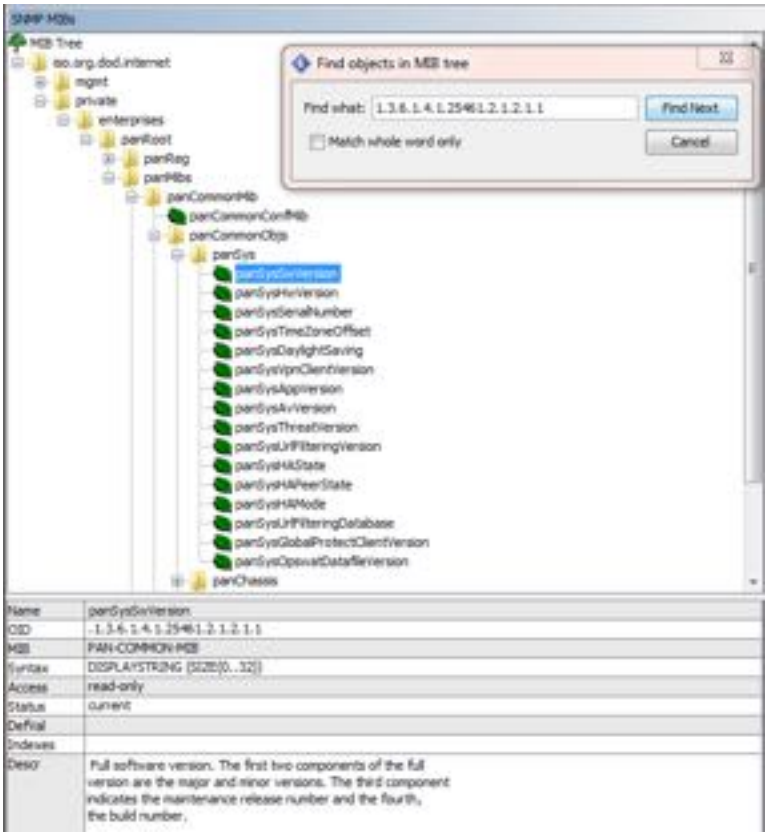
Identify a MIB Containing a Known OID

If you already know the OID for a particular SNMP object (statistic or trap) and want to know the OIDs of similar objects so you can monitor them, you can explore the MIB that contains the known OID.

STEP 1 | Load all the Supported MIBs on page 425 into your SNMP manager.

© 2018 Palo Alto Networks, Inc.

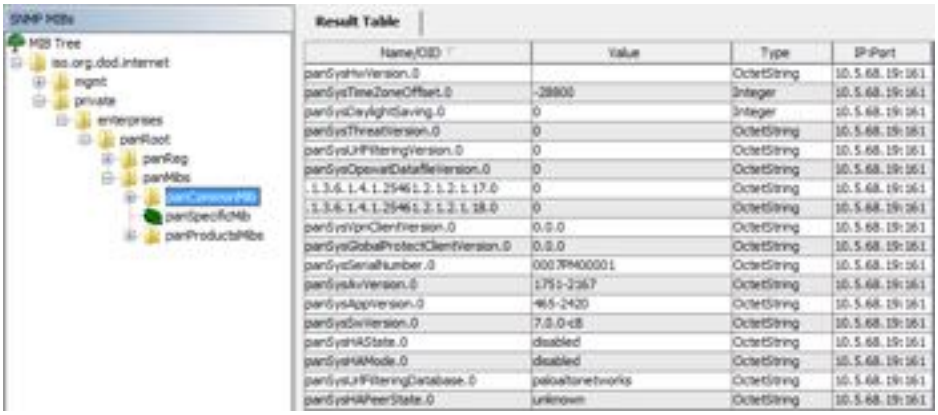
STEP 2 | Search the entire MIB tree for the known OID. The search result displays the MIB path for the OID, as well as information about the OID (for example, name, status, and description). You can then select other OIDs in the same MIB to see information about them.



STEP 3 | (Optional) Walk a MIB on page 420 to display all its objects.

Walk a MIB

If you want to see which SNMP objects (system statistics and traps) are available for monitoring, displaying all the objects of a particular MIB can be useful. To do this, load the Supported MIBs on page 425 into your SNMP manager and perform a *walk* on the desired MIB. To list the traps that Palo Alto Networks firewalls, Panorama, and WF-500 appliance support, walk the panCommonEventEventsV2 MIB. In the following example, walking the PAN-COMMON-MIB.my on page 430 displays the following list of OIDs and their values for certain statistics:



© 2018 Palo Alto Networks, Inc.

Identify the OID for a System Statistic or Trap

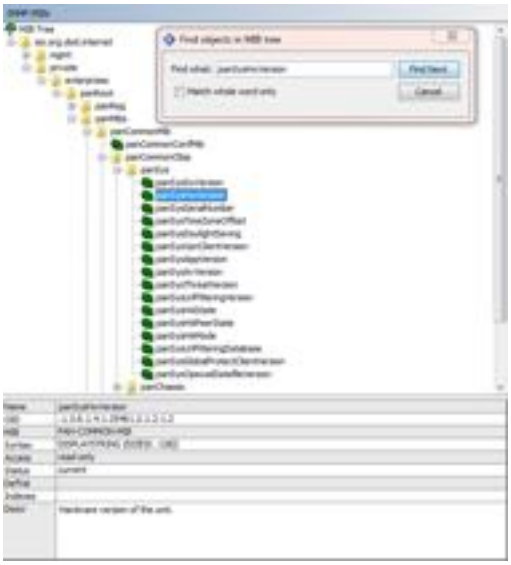
To use an SNMP manager for monitoring Palo Alto Networks firewalls, Panorama, or WF-500 appliances, you must know the OIDs of the system statistics and traps you want to monitor.

STEP 1 | Review the Supported MIBs to determine which one contains the type of statistic you want. For example, the PAN-COMMON-MIB.my contains hardware version information. The panCommonEventEventsV2 MIB contains all the traps that Palo Alto Networks firewalls, Panorama, and WF-500 appliances support.

STEP 2 | Open the MIB in a text editor and perform a keyword search. For example, using **Hardware version** as a search string in PAN-COMMON-MIB identifies the panSysHwVersion object:

```
panSysHwVersion OBJECT-TYPE
 SYNTAX DisplayString (SIZE(0..128))
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION
 "Hardware version of the unit."
 ::= {panSys 2}
```

STEP 3 | In a MIB browser, search the MIB tree for the identified object name to display its OID. For example, the panSysHwVersion object has an OID of 1.3.6.1.4.1.25461.2.1.2.1.2.



Enable SNMP Services for Firewall-Secured Network Elements

If you will use Simple Network Management Protocol (SNMP) to monitor or manage network elements (for example, switches and routers) that are within the security zones of Palo Alto Networks firewalls, you must create a security rule that allows SNMP services for those elements.

You don't need a security rule to enable SNMP monitoring of Palo Alto Networks firewalls, Panorama, or WF-500 appliances. For details, see Monitor Statistics Using SNMP *on page 422.*

STEP 1 | Create an application group.

© 2018 Palo Alto Networks, Inc.

1. Select **Objects** > **Application Group** and click **Add**.
2. Enter a **Name** to identify the application group.
3. Click **Add**, type **snmp**, and select **snmp** and **snmp-trap** from the drop-down.
4. Click **OK** to save the application group.

STEP 2 | Create a security rule to allow SNMP services.

1. Select **Policies** > **Security** and click **Add**.
2. In the **General** tab, enter a **Name** for the rule.
3. In the **Source** and **Destination** tabs, click **Add** and enter a **Source Zone** and a **Destination Zone** for the traffic.
4. In the **Applications** tab, click **Add**, type the name of the applications group you just created, and select it from the drop-down.
5. In the **Actions** tab, verify that the **Action** is set to **Allow**, and then click **OK** and **Commit**.

Monitor Statistics Using SNMP

The statistics that a Simple Network Management Protocol (SNMP) manager collects from Palo Alto Networks firewalls can help you gauge the health of your network (systems and connections), identify resource limitations, and monitor traffic or processing loads. The statistics include information such as interface states (up or down), active user sessions, concurrent sessions, session utilization, temperature, and system uptime.

You can't configure an SNMP manager to control Palo Alto Networks firewalls (using SET messages), only to collect statistics from them (using GET messages). For details on how SNMP is implemented for Palo Alto Networks firewalls, see SNMP Support.

STEP 1 | Configure the SNMP Manager to get statistics from firewalls.

The following steps provide an overview of the tasks you perform on the SNMP manager. For the specific steps, refer to the documentation of your SNMP manager.

1. To enable the SNMP manager to interpret firewall statistics, load the Supported MIBs for Palo Alto Networks firewalls and, if necessary, compile them.
2. For each firewall that the SNMP manager will monitor, define the connection settings (IP address and port) and authentication settings (SNMPv2c community string or SNMPv3 EngineID/username/password) for the firewall.

   All Palo Alto Networks firewalls use port 161.

   The SNMP manager can use the same or different connection and authentication settings for multiple firewalls. The settings must match those you define when you configure SNMP on the firewall (see Step 3). For example, if you use SNMPv2c, the community string you define when configuring the firewall must match the community string you define in the SNMP manager for that firewall.
3. Determine the object identifiers (OIDs) of the statistics you want to monitor. For example, to monitor the session utilization percentage of a firewall, a MIB browser shows that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.1.2.3.1.0 in PAN-COMMON-MIB.my. For details, see Use an SNMP Manager to Explore MIBs and Objects.
4. Configure the SNMP manager to monitor the desired OIDs.

STEP 2 | Enable SNMP traffic on a firewall interface.

This is the interface that will receive statistics requests from the SNMP manager.

© 2018 Palo Alto Networks, Inc.

PAN-OS doesn't synchronize management (MGT) interface settings for firewalls in a high availability (HA) configuration. You must configure the interface for each HA peer.

Perform this step in the firewall web interface.

- To enable SNMP traffic on the MGT interface, select **Device** > **Setup** > **Interfaces**, edit the **Management** interface, select **SNMP**, and then click **OK** and **Commit**.
- To enable SNMP traffic on any other interface, create an interface management profile for SNMP services and assign the profile to the interface that will receive the SNMP requests. The interface type must be Layer 3 Ethernet.

STEP 3 | Configure the firewall to respond to statistics requests from an SNMP manager.

PAN-OS doesn't synchronize SNMP response settings for firewalls in a high availability (HA) configuration. You must configure these settings for each HA peer.

1. Select **Device** > **Setup** > **Operations** and, in the Miscellaneous section, click **SNMP Setup**.
2. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP Support.
   - **V2c**—Enter the **SNMP Community String**, which identifies a community of SNMP managers and monitored devices, and serves as a password to authenticate the community members to each other.

     *As a best practice, don't use the default community string **`public`**; it's well known and therefore not secure.*
   - **V3**—Create at least one SNMP view group and one user. User accounts and views provide authentication, privacy, and access control when firewalls forward traps and SNMP managers get firewall statistics.
     - **Views**—Each view is a paired OID and bitwise mask: the OID specifies a MIB and the mask (in hexadecimal format) specifies which objects are accessible within (include matching) or outside (exclude matching) that MIB. Click **Add** in the first list and enter a **Name** for the group of views. For each view in the group, click **Add** and configure the view **Name**, **OID**, matching **Option** (**include** or **exclude**), and **Mask**.
     - **Users**—Click **Add** in the second list, enter a username under **Users**, select the **View** group from the drop-down, enter the authentication password (**Auth Password**) used to authenticate to the SNMP manager, and enter the privacy password (**Priv Password**) used to encrypt SNMP messages to the SNMP manager.
3. Click **OK** and **Commit**.

STEP 4 | Monitor the firewall statistics in an SNMP manager.

Refer to the documentation of your SNMP manager for details.

When monitoring statistics related to firewall interfaces, you must match the interface indexes in the SNMP manager with interface names in the firewall web interface. For details, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.

Forward Traps to an SNMP Manager

Simple Network Management Protocol (SNMP) traps can alert you to system events (failures or changes in hardware or software of Palo Alto Networks firewalls) or to threats (traffic that matches a firewall security rule) that require immediate attention.

© 2018 Palo Alto Networks, Inc.

To see the list of traps that Palo Alto Networks firewalls support, use your SNMP Manager to access the panCommonEventEventsV2 MIB. For details, see Use an SNMP Manager to Explore MIBs and Objects.

For details on how for Palo Alto Networks firewalls implement SNMP, see SNMP Support.

STEP 1 | Enable the SNMP manager to interpret the traps it receives.

Load the Supported MIBs for Palo Alto Networks firewalls and, if necessary, compile them. For the specific steps, refer to the documentation of your SNMP manager.

STEP 2 | Configure an SNMP Trap server profile.

The profile defines how the firewall accesses the SNMP managers (trap servers). You can define up to four SNMP managers for each profile.

Optionally, configure separate SNMP Trap server profiles for different log types, severity levels, and WildFire verdicts.

1. Log in to the firewall web interface.
2. Select **Device** > **Server Profiles** > **SNMP Trap**.
3. Click **Add** and enter a **Name** for the profile.
4. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available.
5. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP Support.
   - **V2c**—For each server, click **Add** and enter the server **Name**, IP address (**SNMP Manager**), and **Community String**. The community string identifies a community of SNMP managers and monitored devices, and serves as a password to authenticate the community members to each other.

     As a best practice, don't use the default community string ***public****; it's well known and therefore not secure.*
   - **V3**—For each server, click **Add** and enter the server **Name**, IP address (**SNMP Manager**), SNMP **User** account (this must match a username defined in the SNMP manager), **EngineID** used to uniquely identify the firewall (you can leave the field blank to use the firewall serial number), authentication password (**Auth Password**) used to authenticate to the server, and privacy password (**Priv Password**) used to encrypt SNMP messages to the server.
6. Click **OK** to save the server profile.

STEP 3 | Configure log forwarding.

1. Configure the destinations of Traffic, Threat, and WildFire traps:
   1. Create a Log Forwarding profile. For each log type and each severity level or WildFire verdict, select the **SNMP Trap** server profile.
   2. Assign the Log Forwarding profile to policy rules and network zones. The rules and zones will trigger trap generation and forwarding.
2. Configure the destinations for System, Configuration, User-ID, HIP Match, and Correlation logs. For each log (trap) type and severity level, select the **SNMP Trap** server profile.
3. Click **Commit**.

STEP 4 | Monitor the traps in an SNMP manager.

Refer to the documentation of your SNMP manager.

© 2018 Palo Alto Networks, Inc.

When monitoring traps related to firewall interfaces, you must match the interface indexes in the SNMP manager with interface names in the firewall web interface. For details, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.

Supported MIBs

The following table lists the Simple Network Management Protocol (SNMP) management information bases (MIBs) that Palo Alto Networks firewalls, Panorama, and WF-500 appliances support. You must load these MIBs into your SNMP manager to monitor the objects (system statistics and traps) that are defined in the MIBs. For details, see Use an SNMP Manager to Explore MIBs and Objects.

MIB Type	Supported MIBs
Standard—The Internet Engineering Task Force (IETF) maintains most standard MIBs. You can download the MIBs from the IETF website. *Palo Alto Networks firewalls, Panorama, and WF-500 appliances don't support every object (OID) in every one of these MIBs. See the Supported MIBs links for an overview of the supported OIDs.*	MIB-II IF-MIB HOST-RESOURCES-MIB ENTITY-MIB ENTITY-SENSOR-MIB ENTITY-STATE-MIB IEEE 802.3 LAG MIB LLDP-V2-MIB.my BFD-STD-MIB
Enterprise—You can download the enterprise MIBs from the Palo Alto Networks Technical Documentation portal.	PAN-COMMON-MIB.my PAN-GLOBAL-REG-MIB.my PAN-GLOBAL-TC-MIB.my PAN-LC-MIB.my PAN-PRODUCT-MIB.my PAN-ENTITY-EXT-MIB.my PAN-TRAPS.my

MIB-II

MIB-II provides object identifiers (OIDs) for network management protocols in TCP/IP-based networks. Use this MIB to monitor general information about systems and interfaces. For example, you can analyze trends

© 2018 Palo Alto Networks, Inc.

in bandwidth usage by interface type (ifType object) to determine if the firewall needs more interfaces of that type to accommodate spikes in traffic volume.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support only the following object groups:

Object Group	Description
system	Provides system information such as the hardware model, system uptime, FQDN, and physical location.
interfaces	Provides statistics for physical and logical interfaces such as type, current bandwidth (speed), operational status (for example, up or down), and discarded packets. Logical interface support includes VPN tunnels, aggregate groups, Layer 2 subinterfaces, Layer 3 subinterfaces, loopback interfaces, and VLAN interfaces.

RFC 1213 defines this MIB.

IF-MIB

IF-MIB supports interface types (physical and logical) and larger counters (64K) beyond those defined in MIB-II on page 425. Use this MIB to monitor interface statistics in addition to those that MIB-II provides. For example, to monitor the current bandwidth of high-speed interfaces (greater than 2.2Gps) such as the 10G interfaces of the PA-5000 Series firewalls, you must check the ifHighSpeed object in IF-MIB instead of the ifSpeed object in MIB-II. IF-MIB statistics can be useful when evaluating the capacity of your network.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support only the ifXTable in IF-MIB, which provides interface information such as the number of multicast and broadcast packets transmitted and received, whether an interface is in promiscuous mode, and whether an interface has a physical connector.

RFC 2863 defines this MIB.

HOST-RESOURCES-MIB

HOST-RESOURCES-MIB provides information for host computer resources. Use this MIB to monitor CPU and memory usage statistics. For example, checking the current CPU load (hrProcessorLoad object) can help you troubleshoot performance issues on the firewall.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support portions of the following object groups:

Object Group	Description
hrDevice	Provides information such as CPU load, storage capacity, and partition size. The hrProcessorLoad OIDs provide an average of the cores that process packets. For the PA-5060 firewall, which has multiple dataplanes (DPs), the average is of the cores across all the three DPs that process packets.
hrSystem	Provides information such as system uptime, number of current user sessions, and number of current processes.
hrStorage	Provides information such as the amount of used storage.

RFC 2790 defines this MIB.

© 2018 Palo Alto Networks, Inc.

ENTITY-MIB

ENTITY-MIB provides OIDs for multiple logical and physical components. Use this MIB to determine what physical components are loaded on a system (for example, fans and temperature sensors) and see related information such as models and serial numbers. You can also use the index numbers for these components to determine their operational status in the ENTITY-SENSOR-MIB on page 428 and ENTITY-STATE-MIB on page 428.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support only portions of the entPhysicalTable group:

Object	Description
entPhysicalIndex	A single namespace that includes disk slots and disk drives.
entPhysicalDescr	The component description.
entPhysicalVendorType	The sysObjectID (see PAN-PRODUCT-MIB.my on page 431) when it is available (chassis and module objects).
entPhysicalContainedIn	The value of entPhysicalIndex for the component that contains this component.
entPhysicalClass	Chassis (3), container (5) for a slot, power supply (6), fan (7), sensor (8) for each temperature or other environmental, and module (9) for each line card.
entPhysicalParentRelPos	The relative position of this *child* component among its *sibling* components. Sibling components are defined as entPhysicalEntry components that share the same instance values of each of the entPhysicalContainedIn and entPhysicalClass objects.
entPhysicalName	Supported only if the management (MGT) interface allows for naming the line card.
entPhysicalHardwareRev	The vendor-specific hardware revision of the component.
entPhysicalFirwareRev	The vendor-specific firmware revision of the component.
entPhysicalSoftwareRev	The vendor-specific software revision of the component.
entPhysicalSerialNum	The vendor-specific serial number of the component.
entPhysicalMfgName	The name of the manufacturer of the component.
entPhysicalMfgDate	The date when the component was manufactured.
entPhysicalModelName	The disk model number.
entPhysicalAlias	An alias that the network manager specified for the component.
entPhysicalAssetID	A user-assigned asset tracking identifier that the network manager specified for the component.

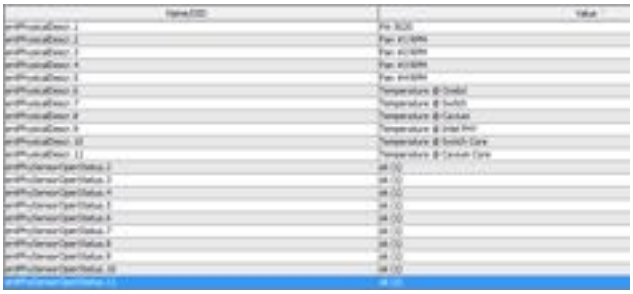
© 2018 Palo Alto Networks, Inc.

Object	Description
entPhysicalIsFRU	Indicates whether the component is a field replaceable unit (FRU).
entPhysicalUris	The Common Language Equipment Identifier (CLEI) number of the component (for example, URN:CLEI:CNME120ARA).

RFC 4133 defines this MIB.

ENTITY-SENSOR-MIB

ENTITY-SENSOR-MIB adds support for physical sensors of networking equipment beyond what ENTITY-MIB on page 427 defines. Use this MIB in tandem with the ENTITY-MIB to monitor the operational status of the physical components of a system (for example, fans and temperature sensors). For example, to troubleshoot issues that might result from environmental conditions, you can map the entity indexes from the ENTITY-MIB (entPhysicalDescr object) to operational status values (entPhysSensorOperStatus object) in the ENTITY-SENSOR-MIB. In the following example, all the fans and temperature sensors for a PA-3020 firewall are working:



The same OID might refer to different sensors on different platforms. Use the ENTITY-MIB for the targeted platform to match the value to the description.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support only portions of the entPhySensorTable group. The supported portions vary by platform and include only thermal (temperature in Celsius) and fan (in RPM) sensors.

RFC 3433 defines the ENTITY-SENSOR-MIB.

ENTITY-STATE-MIB

ENTITY-STATE-MIB provides information about the state of physical components beyond what ENTITY-MIB on page 427 defines, including the administrative and operational state of components in chassis-based platforms. Use this MIB in tandem with the ENTITY-MIB to monitor the operational state of the components of a PA-7000 Series firewall (for example, line cards, fan trays, and power supplies). For example, to troubleshoot log forwarding issues for Threat logs, you can map the log processing card (LPC) indexes from the ENTITY-MIB (entPhysicalDescr object) to operational state values (entStateOper object) in the ENTITY-STATE-MIB. The operational state values use numbers to indicate state: 1 for unknown, 2 for disabled, 3 for enabled, and 4 for testing. The PA-7000 Series firewall is the only Palo Alto Networks firewall that supports this MIB.

RFC 4268 defines the ENTITY-STATE-MIB.

© 2018 Palo Alto Networks, Inc.

IEEE 802.3 LAG MIB

Use the IEEE 802.3 LAG MIB to monitor the status of aggregate groups that have Link Aggregation Control Protocol (ECMP on page 937) enabled. When the firewall logs LACP events, it also generates traps that are useful for troubleshooting. For example, the traps can tell you whether traffic interruptions between the firewall and an LACP peer resulted from lost connectivity or from mismatched interface speed and duplex values.

PAN-OS implements the following SNMP tables for LACP.

The dot3adTablesLastChanged object indicates the time of the most recent change to dot3adAggTable, dot3adAggPortListTable, and dot3adAggPortTable.

Table	Description
Aggregator Configuration Table (dot3adAggTable)	This table contains information about every aggregate group that is associated with a firewall. Each aggregate group has one entry. Some table objects have restrictions, which the dot3adAggIndex object describes. This index is the unique identifier that the local system assigns to the aggregate group. It identifies an aggregate group instance among the subordinate managed objects of the containing object. The identifier is read-only. *The ifTable MIB (a list of interface entries) does not support logical interfaces and therefore does not have an entry for the aggregate group.*
Aggregation Port List Table (dot3adAggPortListTable)	This table lists the ports associated with each aggregate group in a firewall. Each aggregate group has one entry. The dot3adAggPortListPorts attribute lists the complete set of ports associated with an aggregate group. Each bit set in the list represents a port member. For non-chassis platforms, this is a 64-bit value. For chassis platforms, the value is an array of eight 64-bit entries.
Aggregation Port Table (dot3adAggPortTable)	This table contains LACP configuration information about every port associated with an aggregate group in a firewall. Each port has one entry. The table has no entries for ports that are not associated with an aggregate group.
LACP Statistics Table (dot3adAggPortStatsTable)	This table contains link aggregation information about every port associated with an aggregate group in a firewall. Each port has one row. The table has no entries for ports that are not associated with an aggregate group.

The IEEE 802.3 LAG MIB includes the following LACP-related traps:

Trap Name	Description
panLACPLostConnectivityTrap	The peer lost connectivity to the firewall.
panLACPUnresponsiveTrap	The peer does not respond to the firewall.
panLACPNegoFailTrap	LACP negotiation with the peer failed.

© 2018 Palo Alto Networks, Inc.

Trap Name	Description
panLACPSpeedDuplexTrap	The link speed and duplex settings on the firewall and peer do not match.
panLACPLinkDownTrap	An interface in the aggregate group is down.
panLACPLacpDownTrap	An interface was removed from the aggregate group.
panLACPLacpUpTrap	An interface was added to the aggregate group.

For the MIB definitions, refer to IEEE 802.3 LAG MIB.

LLDP-V2-MIB.my

Use the LLDP-V2-MIB to monitor Link Layer Discovery Protocol (LLDP on page 942) events. For example, you can check the lldpV2StatsRxPortFramesDiscardedTotal object to see the number of LLDP frames that were discarded for any reason. The Palo Alto Networks firewall uses LLDP to discover neighboring devices and their capabilities. LLDP makes troubleshooting easier, especially for virtual wire deployments where the ping or traceroute utilities won't detect the firewall.

Palo Alto Networks firewalls support all the LLDP-V2-MIB objects except:

- The following lldpV2Statistics objects:
  - lldpV2StatsRemTablesLastChangeTime
  - lldpV2StatsRemTablesInserts
  - lldpV2StatsRemTablesDeletes
  - lldpV2StatsRemTablesDrops
  - lldpV2StatsRemTablesAgeouts
- The following lldpV2RemoteSystemsData objects:
  - The lldpV2RemOrgDefInfoTable table
  - In the lldpV2RemTable table: lldpV2RemTimeMark

RFC 4957 defines this MIB.

BFD-STD-MIB

Use the Bidirectional Forwarding Detection (BFD) MIB to monitor and receive failure alerts for the bidirectional path between two forwarding engines, such as interfaces, data links, or the actual engines. For example, you can check the bfdSessState object to see the state of a BFD session between forwarding engines. In the Palo Alto Networks implementation, one of the forwarding engines is a firewall interface and the other is an adjacent configured BFD peer.

RFC 7331 defines this MIB.

PAN-COMMON-MIB.my

Use the PAN-COMMON-MIB to monitor the following information for Palo Alto Networks firewalls, Panorama, and WF-500 appliances:

Object Group	Description
panSys	Contains such objects as system software/hardware versions, dynamic content versions, serial number, HA mode/state, and global counters.

© 2018 Palo Alto Networks, Inc.

Object Group	Description
	The global counters include those related to Denial of Service (DoS), IP fragmentation, TCP state, and dropped packets. Tracking these counters enables you to monitor traffic irregularities that result from DoS attacks, system or connection faults, or resource limitations. PAN-COMMON-MIB supports global counters for firewalls but not for Panorama.
panChassis	Chassis type and M-Series appliance mode (Panorama or Log Collector).
panSession	Session utilization information. For example, the total number of active sessions on the firewall or a specific virtual system.
panMgmt	Status of the connection from the firewall to the Panorama management server.
panGlobalProtect	GlobalProtect gateway utilization as a percentage, maximum tunnels allowed, and number of active tunnels.
panLogCollector	Logging statistics for each Log Collector, including logging rate, log quotas, disk usage, retention periods, log redundancy (enabled or disabled), the forwarding status from firewalls to Log Collectors, the forwarding status from Log Collectors to external services, and the status of firewall-to-Log Collector connections.
panDeviceLogging	Logging statistics for each firewall, including logging rate, disk usage, retention periods, the forwarding status from individual firewalls to Panorama and external servers, and the status of firewall-to-Log Collector connections.

PAN-GLOBAL-REG-MIB.my

PAN-GLOBAL-REG-MIB.my contains global, top-level OID definitions for various sub-trees of Palo Alto Networks enterprise MIB modules. This MIB doesn't contain objects for you to monitor; it is required only for referencing by other MIBs.

PAN-GLOBAL-TC-MIB.my

PAN-GLOBAL-TC-MIB.my defines conventions (for example, character length and allowed characters) for the text values of objects in Palo Alto Networks enterprise MIB modules. All Palo Alto Networks products use these conventions. This MIB doesn't contain objects for you to monitor; it is required only for referencing by other MIBs.

PAN-LC-MIB.my

PAN-LC-MIB.my contains definitions of managed objects that Log Collectors (M-Series appliances in Log Collector mode) implement. Use this MIB to monitor the logging rate, log database storage duration (in days), and disk usage (in MB) of each logical disk (up to four) on a Log Collector. For example, you can use this information to determine whether you should add more Log Collectors or forward logs to an external server (for example, a syslog server) for archiving.

PAN-PRODUCT-MIB.my

PAN-PRODUCT-MIB.my defines sysObjectID OIDs for all Palo Alto Networks products. This MIB doesn't contain objects for you to monitor; it is required only for referencing by other MIBs.

© 2018 Palo Alto Networks, Inc.

PAN-ENTITY-EXT-MIB.my

Use PAN-ENTITY-EXT-MIB.my in tandem with the ENTITY-MIB on page 427 to monitor power usage for the physical components of a PA-7000 Series firewall (for example, fan trays, and power supplies), which is the only Palo Alto Networks firewall that supports this MIB. For example, when troubleshooting log forwarding issues, you might want to check the power usage of the log processing cards (LPCs): you can map the LPC indexes from the ENTITY-MIB (entPhysicalDescr object) to values in the PAN-ENTITY-EXT-MIB (panEntryFRUModelPowerUsed object).

PAN-TRAPS.my

Use PAN-TRAPS.my to see a complete listing of all the generated traps and information about them (for example, a description). For a list of traps that Palo Alto Networks firewalls, Panorama, and WF-500 appliances support, refer to the PAN-COMMON-MIB.my on page 430 **panCommonEvents** > **panCommonEventsEvents** > **panCommonEventEventsV2** object.

© 2018 Palo Alto Networks, Inc.

Forward Logs to an HTTP(S) Destination

The firewall and Panorama can forward logs to an HTTP server. You can choose to forward all logs or selectively forward logs to trigger an action on an external HTTP-based service when an event occurs. When forwarding logs to an HTTP server, you can choose the following options:

- Configure the firewall to send an HTTP-based API request directly to a third-party service to trigger an action based on the attributes in a firewall log. You can configure the firewall to work with any HTTP-based service that exposes an API, and modify the URL, HTTP header, parameters, and the payload in the HTTP request to meet your integration needs.
- Tag the source or destination IP address in a log entry automatically and register the IP address and tag mapping to a User-ID agent on the firewall or Panorama, or to a remote User-ID agent so that you can respond to an event and dynamically enforce security policy. To enforce policy, you must Use Dynamic Address Groups in Policy.

STEP 1 | Create an HTTP server profile to forward logs to an HTTP(S) destination.

The HTTP server profile allows you to specify how to access the server and define the format in which to forward logs to the HTTP(S) destination. By default, the firewall uses the management port to forward these logs. You can however, assign a different source interface and IP address in **Device** > **Setup** > **Services** > **Service Route Configuration**.

1. Select **Device** > **Server Profiles** > **HTTP**, add a **Name** for the server profile, and select the **Location**. The profile can be **Shared** across all virtual systems or can belong to a specific virtual system.
2. Click **Add** to provide the details for each server**.** Each profile can have a maximum of 4 servers.
3. Enter a **Name** and IP **Address**.
4. Select the **Protocol** (HTTP or HTTPS). The default **Port** is 80 or 443 respectively; you can modify the port number to match the port on which your HTTP server listens.
5. Select the **HTTP Method** that the third-party service supports—PUT, POST (default), GET and DELETE.
6. Enter the **Username** and **Password** for authenticating to the server, if needed. Click **OK**.
7. Select **Test Server Connection** to verify network connectivity between the firewall and the HTTP(S) server.



STEP 2 | Select the **Payload Format** for the HTTP request.

1. Select the **Log Type** link for each log type for which you want to define the HTTP request format.
2. Select the **Pre-defined Formats** drop-down to view the formats available through content updates, or create a custom format.

   If you create a custom format, the **URI** is the resource endpoint on the HTTP service. The firewall appends the URI to the IP address you defined earlier to construct the URL for the HTTP request. Ensure that the URI and payload format matches the syntax that your third-party vendor requires. You can use any attribute supported on the selected log type within the HTTP Header, Parameter and Value pairs, and in the request payload.

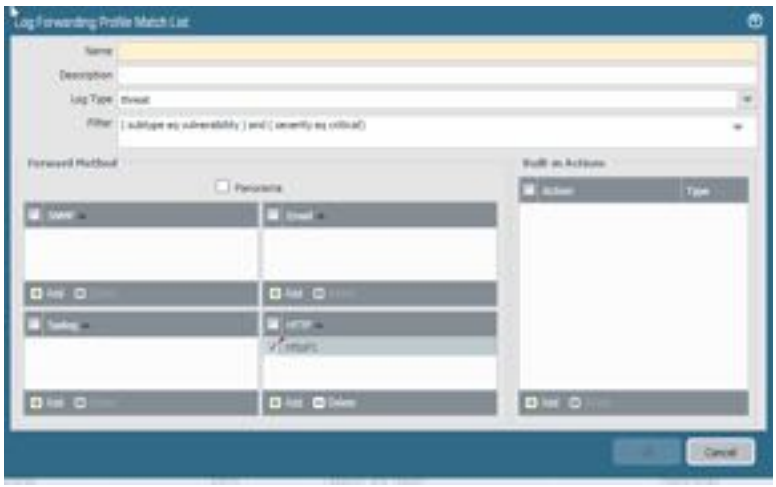
© 2018 Palo Alto Networks, Inc.



3. **Send Test Log** to verify that the HTTP server receives the request. When you interactively send a test log, the firewall uses the format as is and does not replace the variable with a value from a firewall log. If your HTTP server sends a 404 response, provide values for the parameters so that the server can process the request successfully.

STEP 3 | Define the match criteria for when the firewall will forward logs to the HTTP server, and attach the HTTP server profile to use.

1. Select the log types for which you want to trigger a workflow:
   - Add a Log Forwarding Profile (**Objects** > **Log Forwarding Profile**) for logs that pertain to user activity. For example, Traffic, Threat, or Authentication logs.
   - Select **Device** > **Log Settings** for logs that pertain to system events, such as Configuration or System logs.
2. Select the Log Type and use the new **Filter Builder** to define the match criteria.
3. **Add** the HTTP server profile for forwarding logs to the HTTP destination.

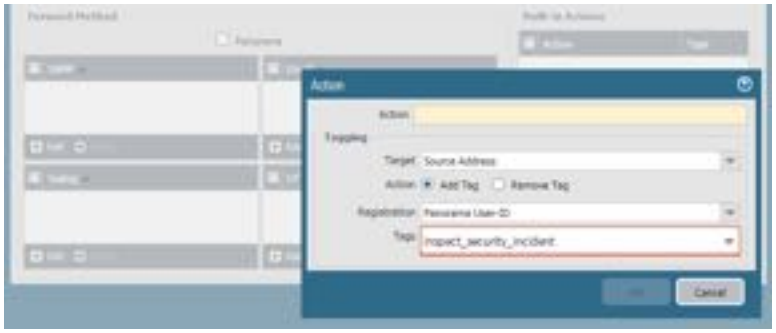


4. Add a tag to the source or destination IP address in the log entry. This capability allows you to use dynamic address groups and security policy rules to limit network access or isolate the IP address until you can triage the affected user device.

   Select **Add** in the Built-in Actions section and select the **Target,** Action: **Add Tag,** and **Registration** to register the tag to the local User-ID on a firewall or to the Panorama that is managing the firewall.

   If you want to register the tag to a remote User-ID agent, see Step 4.

© 2018 Palo Alto Networks, Inc.



STEP 4 | Register or unregister a tag on a source or destination IP address in a log entry to a remote User-ID agent.

1. Select **Device** > **Server Profiles** > **HTTP**, add a **Name** for the server profile, and select the **Location**. The profile can be **Shared** across all virtual systems or can belong to a specific virtual system.
2. Select **Tag Registration** to enable the firewall to register the IP address and tag mapping with the User-ID agent on a remote firewall. With tag registration enabled, you cannot specify the payload format.
3. Add the connection details to access the remote User-ID agent.



4. Select the log type (**Objects** > **Log Forwarding Profile** or **Device** > **Log Settings**) for which you want to add a tag to the source or destination IP address in the log entry.
5. Select **Add** in the Built-in Actions section and **Name** the action. Select the following options to register the tag on the remote User-ID agent:
   - **Target:** Select source or destination IP address.
   - **Action**: Add Tag or Remove Tag.
   - **Registration:** Remote User-ID agent.
   - **HTTP Profile**: Select the profile you created with Tag Registration enabled.
   - **Tag**: Enter a new tag or select from the drop-down.

For dynamic policy enforcement, Use Dynamic Address Groups in Policy.

© 2018 Palo Alto Networks, Inc.

NetFlow Monitoring

NetFlow is an industry-standard protocol that the firewall can use to export statistics about the IP traffic on its interfaces. The firewall exports the statistics as NetFlow fields to a NetFlow collector. The NetFlow collector is a server you use to analyze network traffic for security, administration, accounting and troubleshooting. All Palo Alto Networks firewalls support NetFlow Version 9. The firewalls support only unidirectional NetFlow, not bidirectional. The firewalls perform NetFlow processing on all IP packets on the interfaces and do not support sampled NetFlow. You can export NetFlow records for Layer 3, Layer 2, virtual wire, tap, VLAN, loopback, and tunnel interfaces. For aggregate Ethernet interfaces, you can export records for the aggregate group but not for individual interfaces within the group. To identify firewall interfaces in a NetFlow collector, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors on page 443. The firewalls support standard and enterprise (PAN-OS specific) NetFlow Templates on page 437, which NetFlow collectors use to decipher the NetFlow fields.

- Configure NetFlow Exports on page 436
- NetFlow Templates on page 437

Configure NetFlow Exports

To use a NetFlow collector for analyzing the network traffic on firewall interfaces, perform the following steps to configure NetFlow record exports.

STEP 1 | Create a NetFlow server profile.

The profile defines which NetFlow collectors will receive the exported records and specifies export parameters.

1. Select **Device** > **Server Profiles** > **NetFlow** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Specify the rate at which the firewall refreshes NetFlow Templates in **Minutes** (default is 30) and **Packets** (exported records—default is 20), according to the requirements of your NetFlow collector. The firewall refreshes the templates after either threshold is passed.
4. Specify the **Active Timeout**, which is the frequency in minutes at which the firewall exports records (default is 5).
5. Select **PAN-OS Field Types** if you want the firewall to export App-ID and User-ID fields.
6. **Add** each NetFlow collector (up to two per profile) that will receive records. For each collector, specify the following:
   - **Name** to identify the collector.
   - **NetFlow Server** hostname or IP address.
   - Access **Port** (default 2055).
7. Click **OK** to save the profile.

STEP 2 | Assign the NetFlow server profile to the firewall interfaces that convey the traffic you want to analyze.

In this example, you assign the profile to an existing Ethernet interface.

1. Select **Network** > **Interfaces** > **Ethernet** and click an interface name to edit it.

   You can export NetFlow records for Layer 3, Layer 2, virtual wire, tap, VLAN, loopback, and tunnel interfaces. For aggregate Ethernet interfaces, you can export records for the aggregate group but not for individual interfaces within the group.

2. Select the NetFlow server profile (**NetFlow Profile**) you configured and click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 3 | **(**Required for PA-7000 Series and PA-5200 Series firewalls**)** Configure a service route for the interface that the firewall will use to send NetFlow records.

You cannot use the management (MGT) interface to send NetFlow records from the PA-7000 Series and PA-5200 Series firewalls. For other firewall models, a service route is optional. For all firewalls, the interface that sends NetFlow records does not have to be the same as the interface for which the firewall collects the records.

1. Select **Device** > **Setup** > **Services**.
2. **(Firewall with multiple virtual systems)** Select one of the following:
   - **Global**—Select this option if the service route applies to all virtual systems on the firewall.
   - **Virtual Systems**—Select this option if the service route applies to a specific virtual system. Set the **Location** to the virtual system.
3. Select **Service Route Configuration** and Customize.
4. Select the protocol (**IPv4** or **IPv6**) that the interface uses. You can configure the service route for both protocols if necessary.
5. Click **NetFlow** in the Service column.
6. Select the **Source Interface**.

   Any, Use default, and MGT are not valid interface options for sending NetFlow records from PA-7000 Series or PA-5200 Series firewalls.
7. Select a **Source Address** (IP address).
8. Click **OK** twice to save your changes.

STEP 4 | **Commit** your changes.

STEP 5 | Monitor the firewall traffic in a NetFlow collector.

Refer to your NetFlow collector documentation.



When monitoring statistics, you must match the interface indexes in the NetFlow collector with interface names in the firewall web interface. For details, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.

To troubleshoot NetFlow delivery issues, use the operational CLI command **`debug log-receiver netflow statistics`**.

NetFlow Templates

NetFlow collectors use templates to decipher the fields that the firewall exports. The firewall selects a template based on the type of exported data: IPv4 or IPv6 traffic, with or without NAT, and with standard or enterprise-specific (PAN-OS specific) fields. The firewall periodically refreshes templates to re-evaluate which one to use (in case the type of exported data changes) and to apply any changes to the fields in the selected template. When you Configure NetFlow Exports, set the refresh rate based on a time interval and a number of exported records according to the requirements of your NetFlow collector. The firewall refreshes the templates after either threshold is passed.

The Palo Alto Networks firewall supports the following NetFlow templates:

Template	ID
IPv4 Standard	256
IPv4 Enterprise	257

© 2018 Palo Alto Networks, Inc.

Template	ID
IPv6 Standard	258
IPv6 Enterprise	259
IPv4 with NAT Standard	260
IPv4 with NAT Enterprise	261
IPv6 with NAT Standard	262
IPv6 with NAT Enterprise	263

The following table lists the NetFlow fields that the firewall can send, along with the templates that define them:

Value	Field	Description	Templates
1	IN_BYTES	Incoming counter with length N * 8 bits for the number of bytes associated with an IP flow. By default, N is 4.	All templates
2	IN_PKTS	Incoming counter with length N * 8 bits for the number of packets associated with an IP glow. By default, N is 4.	All templates
4	PROTOCOL	IP protocol byte.	All templates
5	TOS	Type of Service byte setting when entering the ingress interface.	All templates
6	TCP_FLAGS	Total of all the TCP flags in this flow.	All templates
7	L4_SRC_PORT	TCP/UDP source port number (for example, FTP, Telnet, or equivalent).	All templates
8	IPV4_SRC_ADDR	IPv4 source address.	IPv4 standard IPv4 enterprise IPv4 with NAT standard IPv4 with NAT enterprise
10	INPUT_SNMP	Input interface index. The value length is 2 bytes by default, but higher values are possible. For details on how Palo Alto Networks firewalls generate interface indexes, see	All templates

© 2018 Palo Alto Networks, Inc.

Value	Field	Description	Templates
		Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.	
11	L4_DST_PORT	TCP/UDP destination port number (for example, FTP, Telnet, or equivalent).	All templates
12	IPV4_DST_ADDR	IPv4 destination address.	IPv4 standard IPv4 enterprise IPv4 with NAT standard IPv4 with NAT enterprise
14	OUTPUT_SNMP	Output interface index. The value length is 2 bytes by default, but higher values are possible. For details on how Palo Alto Networks firewalls generate interface indexes, see Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.	All templates
21	LAST_SWITCHED	System uptime in milliseconds when the last packet of this flow was switched.	All templates
22	FIRST_SWITCHED	System uptime in milliseconds when the first packet of this flow was switched.	All templates
27	IPV6_SRC_ADDR	IPv6 source address.	IPv6 standard IPv6 enterprise IPv6 with NAT standard IPv6 with NAT enterprise
28	IPV6_DST_ADDR	IPv6 destination address.	IPv6 standard IPv6 enterprise IPv6 with NAT standard IPv6 with NAT enterprise
32	ICMP_TYPE	Internet Control Message Protocol (ICMP) packet type. This is reported as:	All templates

© 2018 Palo Alto Networks, Inc.

Value	Field	Description	Templates
		ICMP Type * 256 + ICMP code	
61	DIRECTION	Flow direction: • 0 = ingress • 1 = egress	All templates
148	flowId	An identifier of a flow that is unique within an observation domain. You can use this information element to distinguish between different flows if flow keys such as IP addresses and port numbers are not reported or are reported in separate records. The flowID corresponds to the session ID field in Traffic and Threat logs.	All templates
233	firewallEvent	Indicates a firewall event: • 0 = Ignore (invalid)—Not used. • 1 = Flow created—The NetFlow data record is for a new flow. • 2 = Flow deleted—The NetFlow data record is for the end of a flow. • 3 = Flow denied—The NetFlow data record indicates a flow that firewall policy denied. • 4 = Flow alert—Not used. • 5 = Flow update—The NetFlow data record is sent for a *long-lasting* flow, which is a flow that lasts longer than the **Active Timeout** period configured in the NetFlow server profile.	All templates
225	postNATSourceIPv4Address	The definition of this information element is identical to that of sourceIPv4Address, except that it reports a modified value that the firewall produced during network address translation after the packet traversed the interface.	IPv4 with NAT standard IPv4 with NAT enterprise
226	postNATDestinationIPv4Address	The definition of this information element is identical to that of destinationIPv4Address, except that it reports a modified value that the firewall produced during network address translation after the packet traversed the interface.	IPv4 with NAT standard IPv4 with NAT enterprise
227	postNAPTSourceTransportPort	The definition of this information element is identical to that of	IPv4 with NAT standard

© 2018 Palo Alto Networks, Inc.

Value	Field	Description	Templates
		sourceTransportPort, except that it reports a modified value that the firewall produced during network address port translation after the packet traversed the interface.	IPv4 with NAT enterprise
228	postNAPTDestinationTransportPort	The definition of this information element is identical to that of destinationTransportPort, except that it reports a modified value that the firewall produced during network address port translation after the packet traversed the interface.	IPv4 with NAT standard IPv4 with NAT enterprise
281	postNATSourceIPv6Address	The definition of this information element is identical to the definition of information element sourceIPv6Address, except that it reports a modified value that the firewall produced during NAT64 network address translation after the packet traversed the interface. See RFC 2460 for the definition of the source address field in the IPv6 header. See RFC 6146 for NAT64 specification.	IPv6 with NAT standard IPv6 with NAT enterprise
282	postNATDestinationIPv6Address	The definition of this information element is identical to the definition of information element destinationIPv6Address, except that it reports a modified value that the firewall produced during NAT64 network address translation after the packet traversed the interface. See RFC 2460 for the definition of the destination address field in the IPv6 header. See RFC 6146 for NAT64 specification.	IPv6 with NAT standard IPv6 with NAT enterprise
346	privateEnterpriseNumber	This is a unique private enterprise number that identifies Palo Alto Networks: 25461.	IPv4 enterprise IPv4 with NAT enterprise IPv6 enterprise IPv6 with NAT enterprise
56701	App-ID	The name of an application that App-ID identified. The name can be up to 32 bytes.	IPv4 enterprise IPv4 with NAT enterprise

© 2018 Palo Alto Networks, Inc.

Value	Field	Description	Templates
			IPv6 enterprise IPv6 with NAT enterprise
56702	User-ID	A username that User-ID identified. The name can be up to 64 bytes.	IPv4 enterprise IPv4 with NAT enterprise IPv6 enterprise IPv6 with NAT enterprise

© 2018 Palo Alto Networks, Inc.

Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors

When you use a NetFlow collector (see NetFlow Monitoring) or SNMP manager (see SNMP Monitoring and Traps) to monitor the Palo Alto Networks firewall, an interface index (SNMP ifindex object) identifies the interface that carried a particular flow (see Interface Indexes in an SNMP Manager). In contrast, the firewall web interface uses interface names as identifiers (for example, ethernet1/1), not indexes. To understand which statistics that you see in a NetFlow collector or SNMP manager apply to which firewall interface, you must be able to match the interface indexes with interface names.

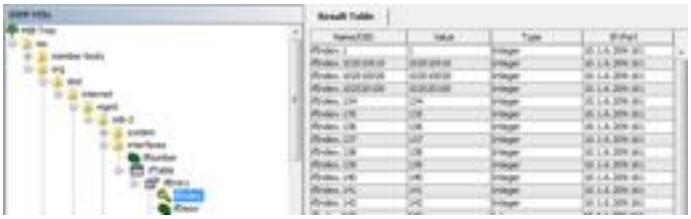


Figure 3: Interface Indexes in an SNMP Manager

You can match the indexes with names by understanding the formulas that the firewall uses to calculate indexes. The formulas vary by platform and interface type: physical or logical.

Physical interface indexes have a range of 1-9999, which the firewall calculates as follows:

Firewall Platform	Calculation	Example Interface Index
VM-Series	Number of management ports + physical port offset • **Number of management ports**—This is a constant of 1. • **Physical port offset**—This is the physical port number.	VM-100 firewall, Eth1/4 = 1 (number of management ports) + 4 (physical port) = **5**
PA-200, PA-500	Number of management ports + physical port offset • **Number of management ports**—This is a constant of 2. • **Physical port offset**—This is the physical port number.	PA-500 Series firewall, Eth1/4 = 2 (number of management ports) + 4 (physical port) = **6**
PA-220, PA-800 Series, PA-3000 Series, PA-5000 Series, Pa-5200 Series	Number of management ports + physical port offset • **Number of management ports**—This is a constant of 5. • **Physical port offset**—This is the physical port number.	PA-5000 Series firewall, Eth1/4 = 5 (number of management ports) + 4 (physical port) = **9**
PA-7000 Series firewalls	(Max. ports * slot) + physical port offset + number of management ports	PA-7000 Series firewall, Eth3/9 =

© 2018 Palo Alto Networks, Inc.

Firewall Platform	Calculation	Example Interface Index
	• **Maximum ports**—This is a constant of 64. • **Slot**—This is the chassis slot number of the network interface card. • **Physical port offset**—This is the physical port number. • **Number of management ports**—This is a constant of 5.	[64 (max. ports) * 3 (slot)] + 9 (physical port) + 5 (number of management ports) = **206**

Logical interface indexes for all platforms are nine-digit numbers that the firewall calculates as follows:

Interface Type	Range	Digit 9	Digits 7-8	Digits 5-6	Digits 1-4	Example Interface Index
Layer 3 subinterface	101010001-199999999	Type: 1	Interface slot: 1-9 (01-09)	Interface port: 1-9 (01-09)	Subinterface: suffix 1-9999 (0001-9999)	Eth1/5.22 = 100000000 (type) + 100000 (slot) + 50000 (port) + 22 (suffix) = **101050022**
Layer 2 subinterface	101010001-199999999	Type: 1	Interface slot: 1-9 (01-09)	Interface port: 1-9 (01-09)	Subinterface: suffix 1-9999 (0001-9999)	Eth2/3.6 = 100000000 (type) + 200000 (slot) + 30000 (port) + 6 (suffix) = **102030006**
Vwire subinterface	101010001-199999999	Type: 1	Interface slot: 1-9 (01-09)	Interface port: 1-9 (01-09)	Subinterface: suffix 1-9999 (0001-9999)	Eth4/2.312 = 100000000 (type) + 400000 (slot) + 20000 (port) + 312 (suffix) = **104020312**
VLAN	200000001-200009999	Type: 2	00	00	VLAN suffix: 1-9999 (0001-9999)	VLAN.55 = 200000000 (type) + 55 (suffix) = **200000055**
Loopback	300000001-300009999	Type: 3	00	00	Loopback suffix: 1-9999 (0001-9999)	Loopback.55 = 300000000 (type) + 55 (suffix) = **300000055**
Tunnel	400000001-400009999	Type: 4	00	00	Tunnel suffix: 1-9999 (0001-9999)	Tunnel.55 = 400000000 (type) + 55 (suffix) = **400000055**
Aggregate group	500010001-500089999	Type: 5	00	AE suffix: 1-8 (01-08)	Subinterface: suffix 1-9999 (0001-9999)	AE5.99 = 500000000 (type) + 50000 (AE Suffix) + 99 (suffix) = **500050099**

User-ID

The user identity, as opposed to an IP address, is an integral component of an effective security infrastructure. Knowing who is using each of the applications on your network, and who may have transmitted a threat or is transferring files, can strengthen security policies and reduce incident response times. User-ID™, a standard feature on the Palo Alto Networks firewall, enables you to leverage user information stored in a wide range of repositories. The following topics provide more details about User-ID and how to configure it:

- User-ID Overview on page 447
- User-ID Concepts on page 449
- Enable User-ID on page 453
- Map Users to Groups on page 456
- Map IP Addresses to Users on page 459
- Enable User- and Group-Based Policy on page 496
- Enable Policy for Users with Multiple Accounts on page 497
- Verify the User-ID Configuration on page 499
- Deploy User-ID in a Large-Scale Network on page 501

© 2018 Palo Alto Networks, Inc.

User-ID Overview

User-ID™ enables you to identify all users on your network using a variety of techniques to ensure that you can identify users in all locations 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 discern what the application is, who the user is, the bandwidth and session consumption, along with the source and destination of the application traffic, as well as any associated threats.
- **Policy control**—Tying user information to Security policy rules improves 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 SaaS applications that enable access to Human Resources services (such as Workday or Service Now) must be available to any known user on your network. However, for more sensitive applications you can reduce your 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 your users do not.
- **Logging, reporting, forensics**—If a security incident occurs, forensics analysis and reporting based on user information rather than just IP addresses provides a more complete picture of the incident. For example, you can use the pre-defined User/Group Activity to see a summary of the web activity of individual users or user groups, or 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 usernames. User-ID provides many mechanisms to collect this User Mapping on page 449 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 Authentication Policy on page 211 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 are safely enabling access to applications for all users, in all locations, all the time.

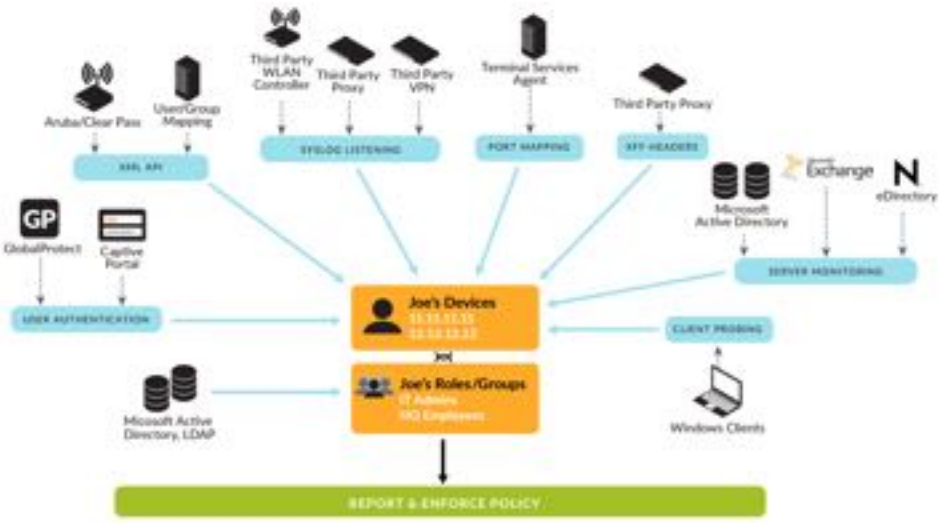


Figure 4: User-ID

© 2018 Palo Alto Networks, Inc.

To enable user- and group-based policy enforcement, the firewall requires a list of all available users and their corresponding group memberships so that you can select groups when defining your policy rules. The firewall collects Group Mapping on page 449 information by connecting directly to your LDAP directory server, or using XML API integration with your directory server.

See User-ID Concepts on page 449 for information on how User-ID works and Enable User-ID on page 453 for instructions on setting up User-ID.



User-ID does not work in environments where the source IP addresses of users are subject to NAT translation before the firewall maps the IP addresses to usernames.

© 2018 Palo Alto Networks, Inc.

User-ID Concepts

- Group Mapping on page 449
- User Mapping on page 449

Group Mapping

To define policy rules based on user or group, first you create an LDAP server profile that defines how the firewall connects and authenticates to your directory server. The firewall supports a variety of directory servers, including Microsoft Active Directory (AD), Novell eDirectory, and Sun ONE Directory Server. The server profile also defines how the firewall searches the directory to retrieve the list of groups and the corresponding list of members. If you are using a directory server that is not natively supported by the firewall, you can integrate the group mapping function using the XML API. You can then create a group mapping configuration to Map Users to Groups on page 456 and Enable User- and Group-Based Policy on page 496.

Defining policy rules based on group membership rather than on individual users simplifies administration because you don't have to update the rules whenever new users are added to a group. When configuring group mapping, you can limit which groups will be available in policy rules. You can specify groups that already exist in your directory service or define custom groups based on LDAP filters. Defining custom groups can be quicker than creating new groups or changing existing ones on an LDAP server, and doesn't require an LDAP administrator to intervene. User-ID maps all the LDAP directory users who match the filter to the custom group. For example, you might want a security policy that allows contractors in the Marketing Department to access social networking sites. If no Active Directory group exists for that department, you can configure an LDAP filter that matches users for whom the LDAP attribute Department is set to Marketing. Log queries and reports that are based on user groups will include custom groups.

User Mapping

Knowing user and groups names is only one piece of the puzzle. The firewall also needs to know which IP addresses map to which users so that security rules can be enforced appropriately. User-ID on page 447 illustrates the different methods that are used to identify users and groups on your network and shows how user mapping and group mapping work together to enable user- and group-based security enforcement and visibility. The following topics describe the different methods of user mapping:

- Server Monitoring on page 449
- Port Mapping on page 450
- Syslog on page 450
- XFF Headers on page 450
- Authentication Policy and Captive Portal on page 450
- GlobalProtect on page 451
- XML API on page 451
- Client Probing on page 452

Server Monitoring

With server monitoring a User-ID agent—either a Windows-based agent running on a domain server in your network, or the integrated PAN-OS User-ID agent running on the firewall—monitors the security event logs for specified Microsoft Exchange Servers, Domain Controllers, or Novell eDirectory servers for login events. For example, in an AD environment, you can configure the User-ID agent to monitor the security logs for Kerberos ticket grants or renewals, Exchange server access (if configured), and file and print service connections. For these events to be recorded in the security log, the AD domain must be configured to log successful account login events. In addition, because users can log in to any of the servers in the domain,

© 2018 Palo Alto Networks, Inc.

you must set up server monitoring for all servers to capture all user login events. See Configure User Mapping Using the Windows User-ID Agent on page 461 or Configure User Mapping Using the PAN-OS Integrated User-ID Agent on page 472 for details.

Port Mapping

In environments with multi-user systems—such as Microsoft Terminal Server or Citrix environments—many users share the same IP address. In this case, the user-to-IP address mapping process requires knowledge of the source port of each client. To perform this type of mapping, you must install the Palo Alto Networks Terminal Services Agent on the Windows/Citrix terminal server itself to intermediate the assignment of source ports to the various user processes. For terminal servers that do not support the Terminal Services agent, such as Linux terminal servers, you can use the XML API to send user mapping information from login and logout events to User-ID. See Configure User Mapping for Terminal Server Users on page 487 for configuration details.

XFF Headers

If you have a proxy server deployed between the users on your network and the firewall, the firewall might see the proxy server IP address as the source IP address in HTTP/HTTPS traffic that the proxy forwards rather than the IP address of the client that requested the content. In many cases, the proxy server adds an X-Forwarded-For (XFF) header to traffic packets that includes the actual IPv4 or IPv6 address of the client that requested the content or from whom the request originated. In such cases, you can configure the firewall to extract the end user IP address from the XFF so that User-ID can map the IP address to a username. This enables you to Use XFF Values for Policies and Logging Source Users so that you can enforce user-based policy to safely enable access to web-based for your users behind a proxy server.

Authentication Policy and Captive Portal

In some cases, the User-ID agent can't map an IP address to a username using server monitoring or other methods—for example, if the user isn't logged in or uses an operating system such as Linux that your domain servers don't support. In other cases, you might want users to authenticate when accessing sensitive applications regardless of which methods the User-ID agent uses to perform user mapping. For all these cases, you can configure Configure Authentication Policy and Map IP Addresses to Usernames Using Captive Portal. Any web traffic (HTTP or HTTPS) that matches an Authentication policy rule prompts the user to authenticate through Captive Portal. You can use the following Captive Portal Authentication Methods:

- Browser challenge—Use Kerberos single sign-on (recommended) or NT LAN Manager (NTLM) authentication if you want to reduce the number of login prompts that users must respond to.
- Web form—Use Multi-Factor Authentication, SAML single sign-on, Kerberos, TACACS+, RADIUS, LDAP, or Local Authentication.
- Client Certificate Authentication.

Syslog

Your environment might have existing network services that authenticate users. These services include wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, and other Network Access Control (NAC) mechanisms. You can configure these services to send syslog messages that contain information about login and logout events and configure the User-ID agent to parse those messages. The User-ID agent parses for login events to map IP addresses to usernames and parses for logout events to delete outdated mappings. Deleting outdated mappings is particularly useful in environments where IP address assignments change often.

Both the PAN-OS integrated User-ID agent and Windows-based User-ID agent use Syslog Parse profiles to parse syslog messages. In environments where services send the messages in different formats, you can create a custom profile for each format and associate multiple profiles with each syslog sender. If you use

© 2018 Palo Alto Networks, Inc.

the PAN-OS integrated User-ID agent, you can also use predefined Syslog Parse profiles that Palo Alto Networks provides through Applications content updates.

Syslog messages must meet the following criteria for a User-ID agent to parse them:

- Each message must be a single-line text string. The allowed delimiters for line breaks are a new line (\n) or a carriage return plus a new line (\r\n).
- The maximum size for individual messages is 2,048 bytes.
- Messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets. A single packet might contain multiple messages.

See Configure User-ID to Monitor Syslog Senders for User Mapping on page 474 for configuration details.

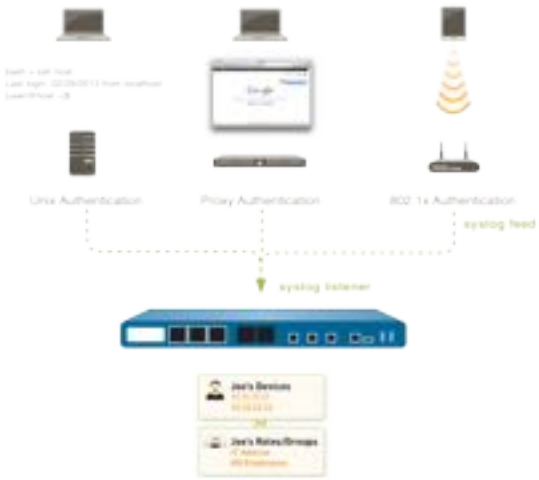


Figure 5: User-ID Integration with Syslog

GlobalProtect

For mobile or roaming users, the GlobalProtect client provides the user mapping information to the firewall directly. In this case, every GlobalProtect user has an agent or app running on the client that requires the user to enter login credentials for VPN access to the firewall. This login information is then added to the User-ID user mapping table on the firewall for visibility and user-based security policy enforcement. Because GlobalProtect users must authenticate to gain access to the network, the IP address-to-username mapping is explicitly known. This is the best solution in sensitive environments where you must be certain of who a user is in order to allow access to an application or service. For more information on setting up GlobalProtect, refer to the GlobalProtect Administrator's Guide.

XML API

Captive Portal and the other standard user mapping methods might not work for certain types of user access. For example, the standard methods cannot add mappings of users connecting from a third-party VPN solution or users connecting to a 802.1x-enabled wireless network. For such cases, you can use the PAN-OS XML API to capture login events and send them to the PAN-OS integrated User-ID agent. See Send User Mappings to User-ID Using the XML API on page 495 for details.

© 2018 Palo Alto Networks, Inc.

Client Probing

In a Microsoft Windows environment, you can configure the User-ID agent to probe client systems using Windows Management Instrumentation (WMI) and/or NetBIOS probing at regular intervals to verify that an existing user mapping is still valid or to obtain the username for an IP address that is not yet mapped.



NetBIOS probing is only supported on the Windows-based User-ID agent; it is not supported on the PAN-OS integrated User-ID agent.

Client probing was designed for legacy networks where most users were on Windows workstations on the internal network, but is not ideal for today's more modern networks that support a roaming and mobile user base on a variety of devices and operating systems. Additionally, client probing can generate a large amount of network traffic (based on the total number of mapped IP addresses) and can pose a security threat when misconfigured. Therefore, client probing is no longer a recommended method for user mapping. Instead collect user mapping information from more isolated and trusted sources, such as domain controllers and through integrations with Syslog on page 450 or the XML API on page 451, which allow you to safely capture user mapping information from any device type or operating system. If you have sensitive applications that require you to know exactly who a user is, configure Authentication Policy and Captive Portal on page 450 to ensure that you are only allowing access to authorized users.



Because WMI probing trusts data reported back from the endpoint, it is not a recommended method of obtaining User-ID information in a high-security network. If you are using the User-ID agent to parse AD security event logs, syslog messages, or the XML API to obtain User-ID mappings, Palo Alto Networks recommends disabling WMI probing.

If you do choose to use WMI probing, do not enable it on external, untrusted interfaces, as this would cause the agent to send WMI probes containing sensitive information such as the username, domain name, and password hash of the User-ID agent service account outside of your network. This information could potentially be exploited by an attacker to penetrate the network to gain further access.

If you do choose to enable probing in your trusted zones, the agent will probe each learned IP address periodically (every 20 minutes by default, but this is configurable) to verify that the same user is still logged in. In addition, when the firewall encounters an IP address for which it has no user mapping, it will send the address to the agent for an immediate probe.

See Configure User Mapping Using the Windows User-ID Agent on page 461 or Configure User Mapping Using the PAN-OS Integrated User-ID Agent on page 472 for details.

© 2018 Palo Alto Networks, Inc.

Enable User-ID

The user identity, as opposed to an IP address, is an integral component of an effective security infrastructure. Knowing who is using each of the applications on your network, and who may have transmitted a threat or is transferring files, can strengthen your security policy and reduce incident response times. User-ID enables you to leverage user information stored in a wide range of repositories for visibility, user- and group-based policy control, and improved logging, reporting, and forensics:

STEP 1 | Enable User-ID on the source zones that contain the users who will send requests that require user-based access controls.



Enable User-ID on trusted zones only. If you enable User-ID and client probing on an external untrusted zone (such as the internet), probes could be sent outside your protected network, resulting in an information disclosure of the User-ID agent service account name, domain name, and encrypted password hash, which could allow an attacker to gain unauthorized access to protected services and applications.

1. Select **Network** > **Zones** and click the Name of the zone.
2. **Enable User Identification** and click **OK**.

STEP 2 | Create a Dedicated Service Account for the User-ID Agent.



As a best practice, create a service account with the minimum set of permissions required to support the User-ID options you enable to reduce your attack surface in the event that the service account is compromised.

This is required if you plan to use the Windows-based User-ID agent or the PAN-OS integrated User-ID agent to monitor domain controllers, Microsoft Exchange servers, or Windows clients for user login and logout events.

STEP 3 | Map Users to Groups.

This enables the firewall to connect to your LDAP directory and retrieve Group Mapping information so that you will be able to select usernames and group names when creating policy.

STEP 4 | Map IP Addresses to Users.



As a best practice, do not enable client probing as a user mapping method on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured.

The way you do this depends on where your users are located and what types of systems they are using, and what systems on your network are collecting login and logout events for your users. You must configure one or more User-ID agents to enable User Mapping:

- Configure User Mapping Using the Windows User-ID Agent.
- Configure User Mapping Using the PAN-OS Integrated User-ID Agent.
- Configure User-ID to Monitor Syslog Senders for User Mapping.
- Configure User Mapping for Terminal Server Users.
- Send User Mappings to User-ID Using the XML API.

STEP 5 | Specify the networks to include and exclude from user mapping.

© 2018 Palo Alto Networks, Inc.



As a best practice, always specify which networks to include and exclude from User-ID. This allows you to ensure that only your trusted assets are probed and that unwanted user mappings are not created unexpectedly.

The way you specify which subnetworks the agent should include or exclude from User-ID depends on whether you are using the Windows-based User-ID agent or the PAN-OSintegrated User-ID agent.

- Specify the subnetworks the Windows User-ID agent should include in or exclude from User-ID.
- Specify the subnetworks the PAN-OS integrated User-ID agent should include in or exclude from user mapping.

STEP 6 | Configure Authentication Policy and Captive Portal.

The firewall uses Captive Portal to authenticate end users when they request services, applications, or URL categories that match Authentication Policy rules. Based on user information collected during authentication, the firewall creates new user mappings or updates existing mappings. The mapping information collected during authentication overrides information collected through other User-ID methods.

1. Configure Captive Portal.
2. Configure Authentication Policy.

STEP 7 | Enable user- and group-based policy enforcement.



Create rules based on group rather than user whenever possible. This prevents you from having to continually update your rules (which requires a commit) whenever your user base changes.

After configuring User-ID, you will be able to choose a username or group name when defining the source or destination of a security rule:

1. Select **Policies** > **Security** and **Add** a new rule or click an existing rule name to edit.
2. Select **User** and specify which users and groups to match in the rule in one of the following ways:
   - If you want to select specific users or groups as matching criteria, click **Add** in the Source User section to display a list of users and groups discovered by the firewall group mapping function. Select the users or groups to add to the rule.
   - If you want to match any user who has or has not authenticated and you don't need to know the specific user or group name, select **known-user** or **unknown** from the drop-down above the Source User list.
3. Configure the rest of the rule as appropriate and then click **OK** to save it. For details on other fields in the security rule, see Set Up a Basic Security Policy.

STEP 8 | Create the Security policy rules to safely enable User-ID within your trusted zones and prevent User-ID traffic from egressing your network.

Follow the Security policy best practices to ensure that the User-ID application (paloalto-userid-agent) is only allowed in the zones where your agents (both your Windows agents and your PAN-OS integrated agents) are monitoring services and distributing mappings to firewalls. Specifically:

- Allow the paloalto-userid-agent application between the zones where your agents reside and the zones where the monitored servers reside (or even better, between the specific systems that host the agent and the monitored servers).
- Allow the paloalto-userid-agent application between the agents and the firewalls that need the user mappings and between firewalls that are redistributing user mappings and the firewalls they are redistributing the information to.

© 2018 Palo Alto Networks, Inc.

- Deny the paloalto-userid-agent application to any external zone, such as your internet zone.

STEP 9 | Configure the firewall to obtain user IP addresses from X-Forwarded-For (XFF) headers.

When the firewall is between the Internet and a proxy server, the IP addresses in the packets that the firewall sees are for the proxy server rather than users. To enable visibility of user IP addresses instead, configure the firewall to use the XFF headers for user mapping. With this option enabled, the firewall matches the IP addresses with usernames referenced in policy to enable control and visibility for the associated users and groups. For details, see Identify Users Connected through a Proxy Server.

1. Select **Device** > **Setup** > **Content-ID** and edit the X-Forwarded-For Headers settings.
2. Select **X-Forwarded-For Header in User-ID**.



Selecting Strip-X-Forwarded-For Header doesn't disable the use of XFF headers for user attribution in policy rules; the firewall zeroes out the XFF value only after using it for user attribution.

3. Click **OK** to save your changes.

STEP 10 | Commit your changes.

Commit your changes to activate them.

STEP 11 | Verify the User-ID Configuration.

After you configure user mapping and group mapping, verify that the configuration works properly and that you can safely enable and monitor user and group access to your applications and services.

© 2018 Palo Alto Networks, Inc.

Map Users to Groups

Defining policy rules based on user group membership rather than individual users simplifies administration because you don't have to update the rules whenever group membership changes. The number of distinct user groups that each firewall or Panorama can reference across all policies varies by model:

- VM-50, VM-100, VM-300, PA-200, PA-220, PA-500, PA-800 Series, PA-3020, and PA-3050 firewalls: 1,000 groups
- VM-500, VM-700, PA-5020, PA-5050, PA-5060, PA-5200 Series, and PA-7000 Series firewalls, and all Panorama models: 10,000 groups

Use the following procedure to enable the firewall to connect to your LDAP directory and retrieve Group Mapping information. You can then Enable User- and Group-Based Policy.



The following are best practices for group mapping in an Active Directory (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 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 and/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/forest. Take steps to ensure unique usernames in separate forests.*
- *If you have Universal Groups, create an LDAP server profile to connect to the Global Catalog server.*

STEP 1 | Add an LDAP server profile.

The profile defines how the firewall connects to the directory servers from which it collects group mapping information.

1. Select **Device** > **Server Profiles** > **LDAP** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** the LDAP servers. You can add up to four servers to the profile but they must be the same **Type**. For each server, enter a **Name** (to identify the server), **LDAP Server** IP address or FQDN, and server **Port** (default 389).
4. Select the server **Type**.

   Based on your selection (such as **active-directory**), the firewall automatically populates the correct LDAP attributes in the group mapping settings. However, if you customized your LDAP schema, you might need to modify the default settings.
5. For the **Base DN**, enter the Distinguished Name (DN) of the LDAP tree location where you want the firewall to start searching for user and group information.
6. For the **Bind DN**, **Password** and **Confirm Password**, enter the authentication credentials for binding to the LDAP tree.

   The **Bind DN** can be a fully qualified LDAP name (such as `cn=administrator,cn=users,dc=acme,dc=local`) or a user principal name (such as `administrator@acme.local`).
7. Enter the **Bind Timeout** and **Search Timeout** in seconds (default is 30 for both).
8. Click **OK** to save the server profile.

© 2018 Palo Alto Networks, Inc.

STEP 2 | Configure the server settings in a group mapping configuration.

1. Select **Device** > **User Identification** > **Group Mapping Settings**.
2. **Add** the group mapping configuration.
3. Enter a unique **Name** to identify the group mapping configuration.
4. Select the LDAP **Server Profile** you just created.
5. (Optional) By default, the **User Domain** field is blank: the firewall automatically detects the domain names for Active Directory (AD) servers. If you enter a value, it overrides any domain names that the firewall retrieves from the LDAP source. Your entry must be the NetBIOS domain name.
6. (Optional) To filter the groups that the firewall tracks for group mapping, in the Group Objects section, enter a **Search Filter** (LDAP query), **Object Class** (group definition), **Group Name**, and **Group Member**.
7. (Optional) To filter the users that the firewall tracks for group mapping, in the User Objects section, enter a **Search Filter** (LDAP query), **Object Class** (user definition), and **User Name**.
8. (Optional) To match User-ID information with email header information identified in the links and attachments of emails forwarded to WildFire™, enter the list of email domains (**Domain List**) in your organization. Use commas to separate multiple domains (up to 256 characters).

   After you click **OK** (later in this procedure), PAN-OS automatically populates the **Mail Attributes** based on the type of LDAP server specified in the **Server Profile**. When a match occurs, the username in the WildFire log email header section will contain a link that opens the **ACC** tab, filtered by user or user group.
9. Make sure the group mapping configuration is **Enabled** (default is enabled).

STEP 3 | Limit which groups will be available in policy rules.

Required only if you want to limit policy rules to specific groups. The combined maximum for the **Group Include List** and **Custom Group** list is 640 entries per group mapping configuration. Each entry can be a single group or a list of groups. By default, if you don't specify groups, all groups are available in policy rules.



*Any custom groups you create will also be available in the Allow List of authentication profiles (*Configure an Authentication Profile and Sequence*).*

1. Add existing groups from the directory service:
   1. Select **Group Include List**.
   2. Select the Available Groups you want to appear in policy rules and add (⊕) them to the Included Groups.
2. If you want to base policy rules on user attributes that don't match existing user groups, create custom groups based on LDAP filters:
   1. Select **Custom Group** and **Add** the group.
   2. Enter a group **Name** that is unique in the group mapping configuration for the current firewall or virtual system.

      If the **Name** has the same value as the Distinguished Name (DN) of an existing AD group domain, the firewall uses the custom group in all references to that name (such as in policies and logs).
   3. Specify an **LDAP Filter** of up to 2,048 UTF-8 characters and click **OK**.

      The firewall doesn't validate LDAP filters, so it's up to you to ensure they are accurate.

      To minimize the performance impact on the LDAP directory server, use only indexed attributes in the filter.
3. Click **OK** to save your changes.

© 2018 Palo Alto Networks, Inc.

A commit is necessary before custom groups will be available in policies and objects.

STEP 4 | **Commit** your changes.

A commit is necessary before you can use custom groups in policies and objects.



After configuring the firewall to retrieve group mapping information from an LDAP server, but before configuring policies based on the groups it retrieves, the best practice is to either wait for the firewall to refresh its group mappings cache or refresh the cache manually. To verify which groups you can currently use in policies, access the firewall CLI *and run the* **`show user group`** *command. To determine when the firewall will next refresh the group mappings cache, run the* **`show user group-mapping statistics`** *command and check the* `Next Action`*. To manually refresh the cache, run the* **`debug user-id refresh group-mapping all`** *command.*

© 2018 Palo Alto Networks, Inc.

Map IP Addresses to Users

User-ID provides many different methods for mapping IP addresses to usernames. Before you begin configuring user mapping, consider where your users are logging in from, what services they are accessing, and what applications and data you need to control access to. This will inform which types of agents or integrations would best allow you to identify your users.

Once you have your plan, you can begin configuring user mapping using one or more of the following methods as needed to enable user-based access and visibility to applications and resources:

- If you have users with client systems that aren't logged in to your domain servers—for example, users running Linux clients that don't log in to the domain—you can Map IP Addresses to Usernames Using Captive Portal. Using Captive Portal in conjunction with Authentication Policy also ensures that all users authenticate to access your most sensitive applications and data.
- To map users as they log in to your Exchange servers, domain controllers, eDirectory servers, or Windows clients you must configure a User-ID agent:
  - Configure User Mapping Using the PAN-OS Integrated User-ID Agent
  - Configure User Mapping Using the Windows User-ID Agent
- If you have clients running multi-user systems in a Windows environment, such as Microsoft Terminal Server or Citrix Metaframe Presentation Server or XenApp, Configure the Palo Alto Networks Terminal Services Agent for User Mapping. For a multi-user system that doesn't run on Windows, you can Retrieve User Mappings from a Terminal Server Using the PAN-OS XML API.
- To obtain user mappings from existing network services that authenticate users—such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms—Configure User-ID to Monitor Syslog Senders for User Mapping.

  While you can configure either the Windows agent or the PAN-OS integrated User-ID agent on the firewall to listen for authentication syslog messages from the network services, because only the PAN-OS integrated agent supports syslog listening over TLS, it is the preferred configuration.
- For other clients that you can't map using the other methods, you can Send User Mappings to User-ID Using the XML API.
- A large-scale network can have hundreds of information sources that firewalls query for user and group mapping and can have numerous firewalls that enforce policies based on the mapping information. You can simplify User-ID administration for such a network by aggregating the mapping information before the User-ID agents collect it. You can also reduce the resources that the firewalls and information sources use in the querying process by configuring some firewalls to redistribute the mapping information. For details, see Deploy User-ID in a Large-Scale Network.

Create a Dedicated Service Account for the User-ID Agent

To use either the Windows-based User-ID agent or the PAN-OS integrated User-ID agent to map users as they log in to your Exchange servers, domain controllers, eDirectory servers, or Windows clients, create a dedicated service account for the User-ID agent on a domain controller in each domain that the agent will monitor.

The required permissions for the service account depend on the user mapping methods and settings you plan to use. For example, if you are using the PAN-OS integrated User-ID agent, the service account requires Server Operator privileges. If you are using the Windows-based User-ID agent, the service account does not require Server Operator privileges. To reduce the risk of compromising the User-ID service account, always configure the account with the minimum set of permissions necessary for the agent to function properly.

© 2018 Palo Alto Networks, Inc.

User-ID provides many methods for safely collecting user mapping information. Some legacy features, which were designed for environments that only required mapping of users on Windows desktops attached to the local network, require privileged service accounts. If the privileged service account is compromised, this would open your network to attack. As a best practice, avoid using legacy features, such as client probing, NTLM authentication, and session monitoring, that require privileges that would pose a threat if compromised. The following workflow details all required privileges and provides guidance for the User-ID features which require privileges that could pose a threat so that you can decide how to best identify users without compromising your overall security posture.

STEP 1 | Create an AD account for the User-ID agent.

You must create a service account in each domain the agent will monitor.

1. Log in to the domain controller.
2. Right-click the Windows icon ( ), **Search** for **`Active Directory Users and Computers`**, and launch the application.
3. In the navigation pane, open the domain tree, right-click **Managed Service Accounts** and select **New** > **User**.
4. Enter the **First Name**, **Last Name**, and **User logon name** of the user and click **Next**.
5. Enter the **Password** and **Confirm Password**, then click **Next** and **Finish**.

STEP 2 | Add the account to the Builtin groups that have privileges for accessing the services and hosts the User-ID agent will monitor.

1. Right-click the service account you just added and **Add to a group**.
2. **Enter the object names to select** as follows to assign the account to groups. Separate each entry with a semicolon.
   - **Event Log Readers** or a custom group that has privileges for reading Security log events. These privileges are required if the User-ID agent will collect mapping information by monitoring Security logs.
   - (**PAN-OS integrated agent only**) **Distributed COM Users** group, which has privileges for launching, activating, and using Distributed Component Object Model (DCOM) objects.
   - (**PAN-OS integrated agent only - Not recommended**) **Server Operators** group, which has privileges for opening sessions. The agent only requires these privileges if you plan to configure it to refresh existing mapping information by monitoring user sessions.

     Because this group also has privileges for shutting down and restarting servers, only assign the account to this group if monitoring user sessions is very important.
   - (**PAN-OS integrated agent only**) If you plan to configure NTLM authentication for Captive Portal, the firewall where you've configured the agent will need to join the domain. To enable this, enter the name of a group that has administrative privileges to join the domain, write to the validated service principal name, and create a computer object within the *computers* organization unit (ou=computers).

     The PAN-OS integrated agent requires privileged operations to join the domain, which poses a security threat if the account is compromised. Consider configuring Kerberos *single sign-on (SSO) or* SAML *SSO authentication for Captive Portal instead of NTLM. Kerberos and SAML are stronger, more secure authentication methods and do not require the firewall to join the domain.*

     For a firewall with multiple virtual systems, only vsys1 can join the domain because of AD restrictions on virtual systems running on the same host.
3. **Check Names** to validate your entries and click **OK** twice.

© 2018 Palo Alto Networks, Inc.

STEP 3 | If you plan to use WMI probing, enable the account to read the CIMV2 namespace and assign the required permissions on the client systems to be probed.



Do not enable client probing on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured. Instead collect user mapping information from more isolated and trusted sources, such as domain controllers and through integrations with Syslog or the XML API, which have the added benefit of allowing you to safely capture user mapping information from any device type or operating system, instead of just Windows clients.

Perform this task on each client system that the User-ID agent will probe for user mapping information:

1. Right-click the Windows icon (), **Search** for **`wmimgmt.msc`**, and launch the WMI Management Console.
2. In the console tree, right-click **WMI Control** and select **Properties**.
3. Select **Security**, select **Root** > **CIMV2**, and click **Security**.
4. **Add** the name of the service account you created, **Check Names** to verify your entry, and click **OK**.



You might have to change the Locations or click Advanced to query for account names. See the dialog help for details.

5. In the Permissions for **`<Username>`** section, **Allow** the **Enable Account** and **Remote Enable** permissions.
6. Click **OK** twice.
7. Use the Local Users and Groups MMC snap-in (lusrmgr.msc) to add the service account to the local Distributed Component Object Model (DCOM) Users and Remote Desktop Users groups on the system that will be probed.

STEP 4 | Turn off account privileges that are not necessary.

By ensuring that the User-ID service account has the minimum set of account privileges, you can reduce the attack surface should the account be compromised.

To ensure that the User-ID account has the minimum privileges necessary, deny the following privileges on the account:

- **Deny interactive logon for the User-ID service account**—While the User-ID service account does need permission to read and parse Active Directory security event logs, it does not require the ability to logon to servers or domain systems interactively. You can restrict this privilege using Group Policies or by using a Managed Service account (refer to Microsoft TechNet for more information).
- **Deny remote access for the User-ID service account**—This prevents an attacker from using the account to access your network from the outside the network.

STEP 5 | Next steps...

You are now ready to:

- Configure User Mapping Using the Windows User-ID Agent.
- Configure User Mapping Using the PAN-OS Integrated User-ID Agent.

Configure User Mapping Using the Windows User-ID Agent

In most cases, the majority of your network users will have logins to your monitored domain services. For these users, the Palo Alto Networks User-ID agent monitors the servers for login events and performs the IP address to username mapping. The way you configure the User-ID agent depends on the size of your environment and the location of your domain servers. As a best practice, locate your User-ID agents near the servers it will monitor (that is, the monitored servers and the Windows User-ID agent should not be

© 2018 Palo Alto Networks, Inc.

across a WAN link from each other). This is because most of the traffic for user mapping occurs between the agent and the monitored server, with only a small amount of traffic—the delta of user mappings since the last update—from the agent to the firewall.

The following topics describe how to install and configure the User-ID Agent and how to configure the firewall to retrieve user mapping information from the agent:

- Install the Windows-Based User-ID Agent on page 462
- Configure the Windows-Based User-ID Agent for User Mapping on page 469

Install the Windows-Based User-ID Agent

The following procedure shows how to install the User-ID agent on a member server in the domain and set up the service account with the required permissions. If you are upgrading, the installer will automatically remove the older version, however, it is a good idea to back up the config.xml file before running the installer.

For information about the system requirements for installing the Windows-based User-ID agent and for information on supported server OS versions, refer to the Palo Alto Networks Compatibility Matrix.

STEP 1 | Create a dedicated Active Directory service account for the User-ID agent to access the services and hosts it will monitor to collect user mappings.

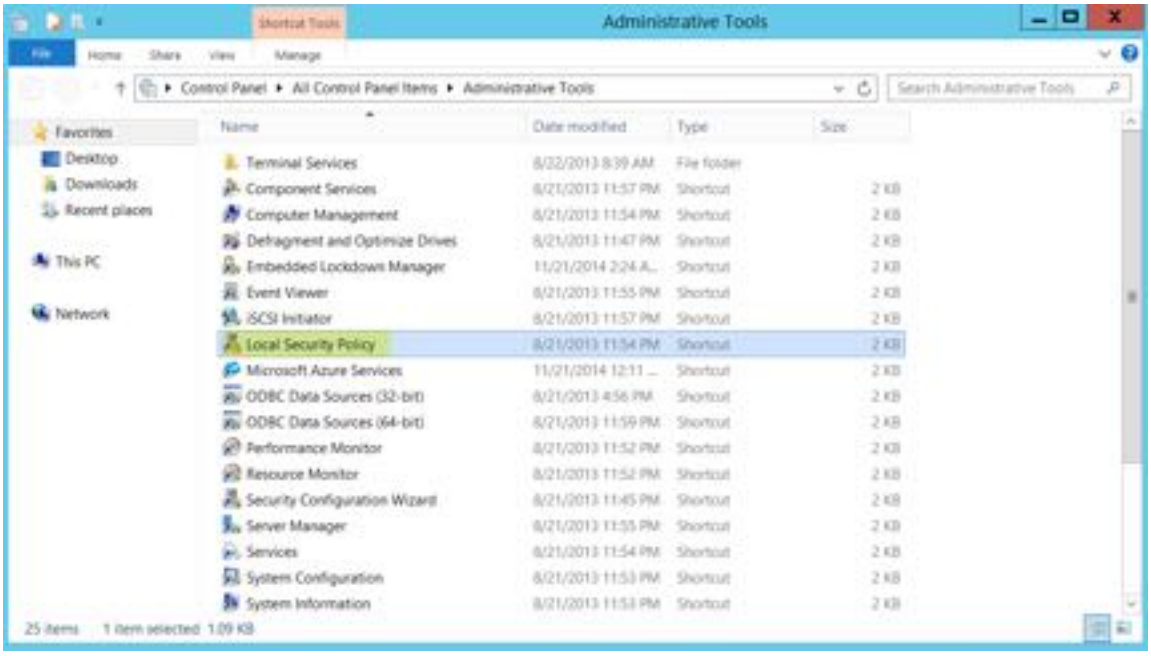
Create a Dedicated Service Account for the User-ID Agent and grant the necessary permissions for the Windows User-ID agent.

1. Enable the service account to log on as a service by configuring either local or group policy.
    1. To configure the group policy if you are installing Windows-based User-ID agents on multiple servers, select **Group Policy Management** > **Default Domain Policy** > **Computer Configuration** > **Policies** > **Windows Settings** > **Security Settings** > **Local Policies** > **User Rights Assignment** for the Windows server that is the agent host.
    2. Right-click **Log on as a service**, then select **Properties**.
    3. Add the service account username or builtin group (Administrators have this privilege by default).

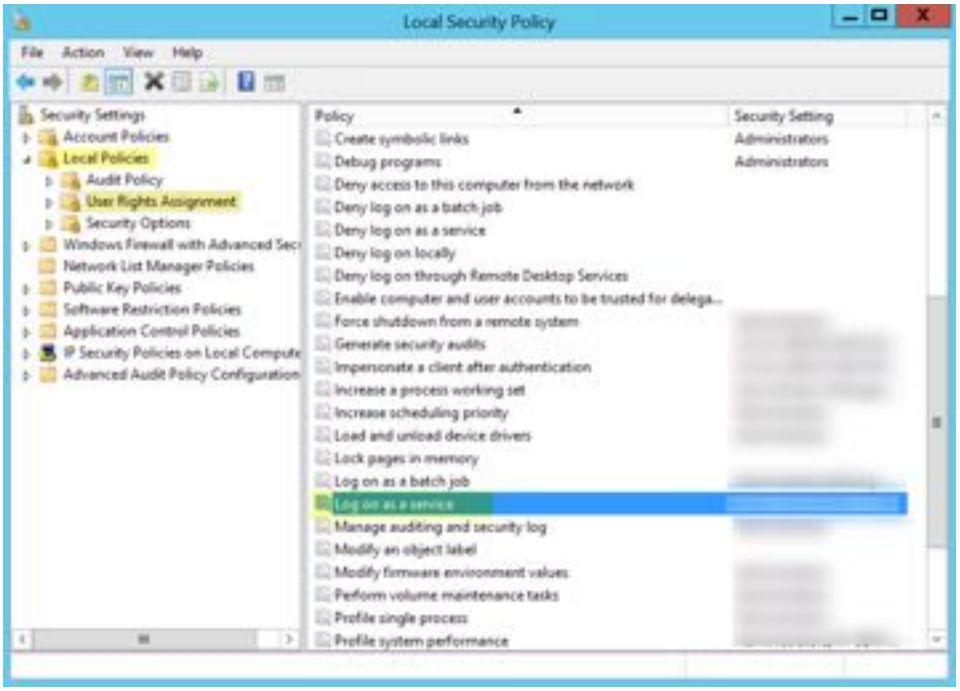
    The permission to log on as a service is only needed locally on the Windows server that is the agent host. If you are using only one User-ID agent, you can grant the permissions locally on the agent host using the following instructions.

    1. To assign permissions locally, select **Control Panel** > **Administrative Tools** > **Local Security Policy**.

© 2018 Palo Alto Networks, Inc.



2. Select **Local Policies** > **User Rights Assignment** > **Log on as a service**.



3. **Add User or Group...** to add the service account.

© 2018 Palo Alto Networks, Inc.



4. Enter the service account name in **`domain\username`** format in the **Enter the object names to select** entry field and click **OK**.



To confirm the service account name is valid, **Check Names**.

2. If you want to use server monitoring to identify users, add the service account to the Event Log Reader builtin group to enable privileges for reading the security log events.
   1. On the domain controller or Exchange server that contains the logs you want the User-ID agent to read, or on the member server that receives events from Windows log forwarding, run the MMC and launch the Active Directory Users and Computers snap-in.
   2. Navigate to the Builtin folder for the domain, right-click the **Event Log Reader** group and select **Add to Group** to open the properties dialog.
   3. Click **Add** and enter the name of the service account that you configured the User-ID service to use and then click **Check Names** to validate that you have the proper object name.
   4. Click **OK** twice to save the settings.
   5. Confirm that the builtin Event Log Reader group lists the service account as a member.

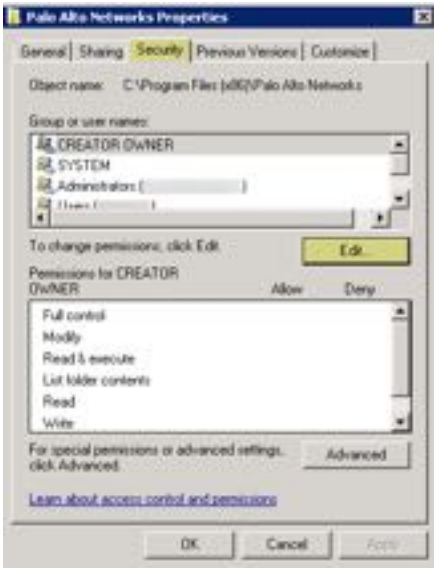
© 2018 Palo Alto Networks, Inc.



3. Assign account permissions to the installation folder to allow the service account to access the agent's installation folder to read the configuration and write logs.

   You only need to perform this step if the service account you configured for the User-ID agent is not either a domain administrator or a local administrator on the User-ID agent server host.

   1. From the Windows Explorer, navigate to **`C:\Program Files(x86)\Palo Alto Networks`** for 32-bit systems, right-click the folder, and select **Properties**.
   2. On the **Security** tab, click **Edit**.



   3. **Add** the User-ID agent service account and assign it permissions to **Modify**, **Read & execute**, **List folder contents**, **Read**, and **Write**, and then click **OK** to save the account settings.

© 2018 Palo Alto Networks, Inc.



If you want to allow the service account to access the User-ID agent's registry keys, Allow the Full Control permission.

4. Give the service account permissions to the User-ID Agent registry sub-tree:
   1. Run **regedt32** and navigate to the Palo Alto Networks sub-tree in one of the following locations:
      - **32-bit systems**—`HKEY_LOCAL_MACHINE\Software\ Palo Alto Networks`
      - **64-bit systems**—`HKEY_LOCAL_MACHINE\Software\ WOW6432Node\Palo Alto Networks`
   2. Right-click the Palo Alto Networks node and select **Permissions**.
   3. Assign the User-ID service account **Full Control** and then click **OK** to save the setting.

STEP 2 | Decide where to install the User-ID agent.

The User-ID agent queries the Domain Controller and Exchange server logs using Microsoft Remote Procedure Calls (MSRPCs). During the initial connection, the agent transfers the most recent 50,000 events from the log to map users. On each subsequent connection, the agent transfers events with a timestamp later than the last communication with the domain controller. Therefore, always install one or more User-ID agents at each site that has servers to be monitored.

- You must install the User-ID agent on a system running one of the supported OS versions: see "Operating System (OS) Compatibility User-ID Agent" in the Compatibility Matrix.
- Make sure the system that will host the User-ID agent is a member of the same domain as the servers it will monitor.
- As a best practice, install the User-ID agent close to the servers it will be monitoring (there is more traffic between the User-ID agent and the monitored servers than there is between the User-ID agent and the firewall, so locating the agent close to the monitored servers optimizes bandwidth usage).
- To ensure the most comprehensive mapping of users, you must monitor all domain controllers that process authentication for users you want to map. You might need to install multiple User-ID agents to efficiently monitor all of your resources.
- If you are using the User-ID agent for credential detection, you must install it on the read-only domain controller (RODC). As a best practice deploy a separate agent for this purpose. Do not use the User-ID agent installed on the RODC to map IP addresses to users. The User-ID agent installer for credential detection is named UaCredInstall64-x.x.x.msi.

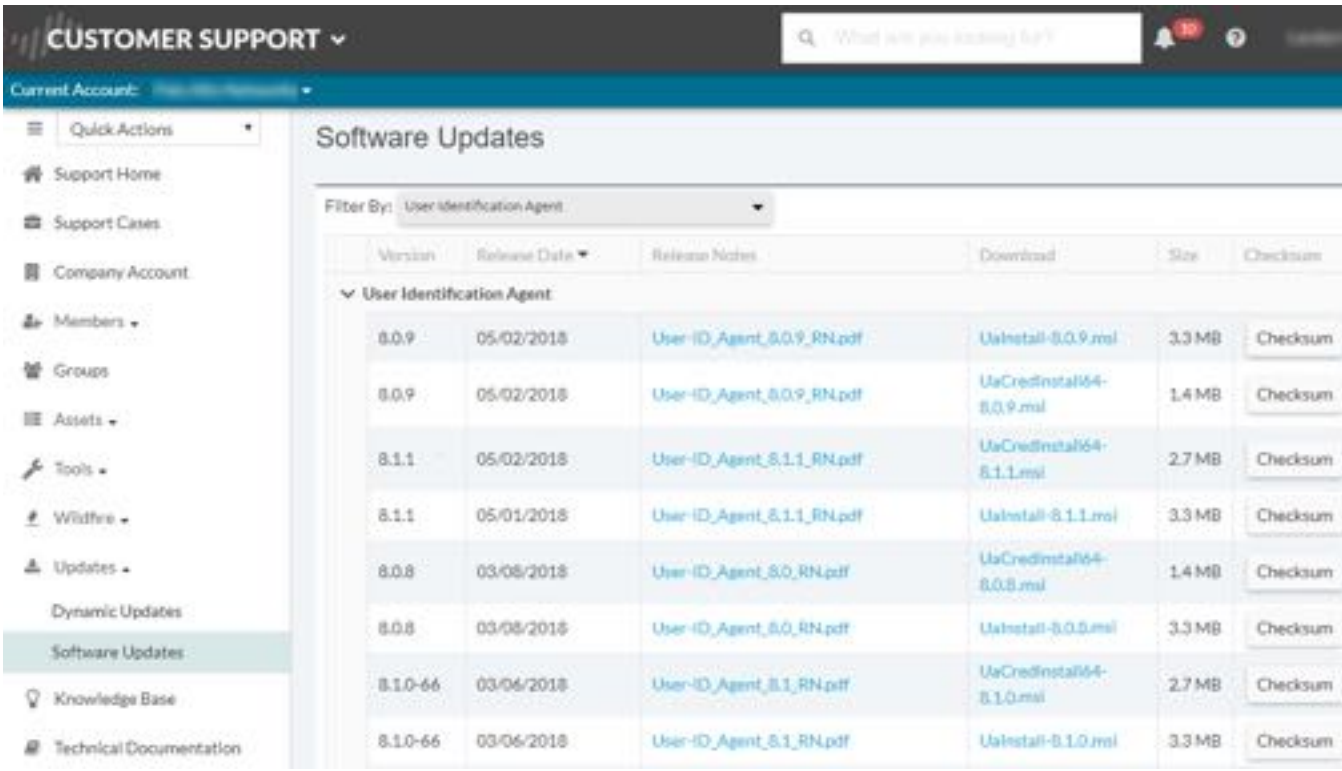
STEP 3 | Download the User-ID agent installer.

Install the User-ID agent version that is the same as the PAN-OS version running on the firewalls. If there is not a User-ID agent version that matches the PAN-OS version, install the latest version that is closest to the PAN-OS version.

© 2018 Palo Alto Networks, Inc.

1. Log in to the Palo Alto Networks Customer Support Portal.
2. Select **Updates** > **Software Updates**.
3. Set **Filter By** to **User Identification Agent** and select the version of the User-ID agent you want to install from the corresponding Download column. For example, to download the 8.0.9 version of the User-ID agent, select **UaInstall-8.0.9.msi**.

   If you are using the User-ID agent for credential detection, make sure you download the `UaCredInstall64-x.x.x.msi` file instead of the regular User-ID installation file, which is named `UaInstall-x.x.x.msi`.
4. Save the `UaInstall.x64-x.x.x-xx.msi` or `UaInstall-x.x.x-xx.msi` file (be sure to select the appropriate version based on whether the Windows system is running a 32-bit OS or a 64-bit OS) on the systems where you plan to install the agent.



STEP 4 | Run the installer as an administrator.

1. Open the Windows **Start** menu, right-click the **Command Prompt** program, and select **Run as administrator**.
2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop, enter the following:

```
C:\Users\administrator.acme>cd Desktop
C:\Users\administrator.acme\Desktop>UaInstall-8.0.9.msi
```

3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to **`C:\Program Files(x86)\Palo Alto Networks`** for 32-bit systems or **`C:\Program Files\Palo Alto Networks`** for 64-bit systems, but you can **Browse** to a different location.
4. When the installation completes, **Close** the setup window.

STEP 5 | Launch the User-ID Agent application as an administrator.

© 2018 Palo Alto Networks, Inc.

Open the Windows **Start** menu, right-click the **User-ID Agent** program, and select **Run as administrator**.

You must run the User-ID Agent application as an administrator to install the application, commit configuration changes, or uninstall the application.

STEP 6 | (Optional) Change the service account that the User-ID agent uses to log in.

By default, the agent uses the administrator account used to install the .msi file. To change the account to a restricted account:

1. Select **User Identification** > **Setup** and click **Edit**.
2. Select the **Authentication** tab and enter the service account name that you want the User-ID agent to use in the **User name for Active Directory** field.
3. Enter the **Password** for the specified account.
4. **Commit** the changes to the User-ID agent configuration to restart the service using the service account credentials.

STEP 7 | (Optional) Assign your own certificates for mutual authentication between the Windows User-ID agent and the firewall.

1. Obtain your certificate for the Windows User-ID agent using one of the following methods. Upload the server certificate in Privacy Enhanced Mail (PEM) format and the server certificate's encrypted key.
   - Generate a Certificate and export it for upload to the Windows User-ID agent.
   - Export a certificate from your enterprise certificate authority (CA) and the upload it to the Windows User-ID agent.
2. Add a server certificate to Windows User-ID agent.
   1. On the Windows User-ID agent, select **Server Certificate** and click **Add**.
   2. Enter the path and name of the certificate file received from the CA or browse to the certificate file.
   3. Enter the private key passphrase.
   4. Click **OK** and then **Commit**.
3. Upload a certificate to the firewall to validate the Windows User-ID agent's identity.
4. Configure the certificate profile for the client device (firewall or Panorama).
   1. Select **Device** > **Certificate Management** > **Certificate Profile**.
   2. Configure a Certificate Profile.

   You can only assign one certificate profile for Windows User-ID agents and Terminal Services (TS) agents. Therefore, your certificate profile must include all certificate authorities that issued certificates uploaded to connected User-ID and TS agents.
5. Assign the certificate profile on the firewall.
   1. Select **Device** > **User Identification** > **Connection Security** and click the edit button.
   2. Select the **User-ID Certificate Profile** you configured in the previous step.
   3. Click **OK**.
6. **Commit** your changes.

STEP 8 | Configure Credential Detection with the Windows-based User-ID Agent.

To use the Windows-based User-ID agent to detect credential submissions and Prevent Credential Phishing, you must install the User-ID credential service on the Windows-based User-ID agent. You can only install this add-on on a read-only domain controller (RODC).

© 2018 Palo Alto Networks, Inc.

Configure the Windows-Based User-ID Agent for User Mapping

The Palo Alto Networks User-ID agent is a Windows service that connects to servers on your network—for example, Active Directory servers, Microsoft Exchange servers, and Novell eDirectory servers—and monitors the logs for login events. The agent uses this information to map IP addresses to usernames. Palo Alto Networks firewalls connect to the User-ID agent to retrieve this user mapping information, enabling visibility into user activity by username rather than IP address and enables user- and group-based security enforcement.



For information about the server OS versions supported by the User-ID agent, refer to "Operating System (OS) Compatibility User-ID Agent" in the User-ID Agent Release Notes*.*

STEP 1 | Define the servers the User-ID agent will monitor to collect IP address to user mapping information.

The User-ID agent can monitor up to 100 servers, of which up to 50 can be syslog senders.



To collect all of the required mappings, the User-ID agent must connect to all servers that your users log in to in order to monitor the security log files on all servers that contain login events.

1. Open the Windows **Start** menu and select **User-ID Agent**.
2. Select **User Identification** > **Discovery**.
3. In the **Servers** section of the screen, click **Add**.
4. Enter a **Name** and **Server Address** for the server to be monitored. The network address can be a FQDN or an IP address.
5. Select the **Server Type** (**Microsoft Active Directory**, **Microsoft Exchange**, **Novell eDirectory**, or **Syslog Sender**) and then click **OK** to save the server entry. Repeat this step for each server to be monitored.
6. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Auto Discover**.

   Auto-discovery locates domain controllers in the local domain only; you must manually add Exchange servers, eDirectory servers, and syslog senders.
7. (Optional) To tune the frequency at which the firewall polls configured servers for mapping information, select **User Identification** > **Setup** and **Edit** the Setup section. On the **Server Monitor** tab, modify the value in the **Server Log Monitor Frequency (seconds)** field. Increase the value in this field to 5 seconds in environments with older Domain Controllers or high-latency links.

   Ensure that the Enable Server Session Read setting is not selected. This setting requires that the User-ID agent have an Active Directory account with Server Operator privileges so that it can read all user sessions. Instead, use a syslog or XML API integration to monitor sources that capture login and logout events for all device types and operating systems (instead of just Windows), such as wireless controllers and Network Access Controllers (NACs).
8. Click **OK** to save the settings.

STEP 2 | Specify the subnetworks the Windows User-ID agent should include in or exclude from User-ID.

By default, the User-ID maps all users accessing the servers you are monitoring.

© 2018 Palo Alto Networks, Inc.

As a best practice, always specify which networks to include and exclude from User-ID to ensure that the agent is only communicating with internal resources and to prevent unauthorized users from being mapped. You should only enable User-ID on the subnetworks where users internal to your organization are logging in.

1. Select **User Identification** > **Discovery**.
2. **Add** an entry to the Include/Exclude list of configured networks and enter a **Name** for the entry and enter the IP address range of the subnetwork in as the **Network Address**.
3. Select whether to include or exclude the network:
   - **Include specified network**—Select this option if you want to limit user mapping to users logged in to the specified subnetwork only. For example, if you include 10.0.0.0/8, the agent maps the users on that subnetwork and excludes all others. If you want the agent to map users in other subnetworks, you must repeat these steps to add additional networks to the list.
   - **Exclude specified network**—Select this option only if you want the agent to exclude a subset of the subnetworks you added for inclusion. For example, if you include 10.0.0.0/8 and exclude 10.2.50.0/22, the agent will map users on all the subnetworks of 10.0.0.0/8 except 10.2.50.0/22, and will exclude all subnetworks outside of 10.0.0.0/8.

     If you add subnetworks for exclusion without adding any for inclusion, the agent will not perform user mapping in any subnetwork.
4. Click **OK**.

STEP 3 | (Optional) If you configured the agent to connect to a Novell eDirectory server, you must specify how the agent should search the directory.

1. Select **User Identification** > **Setup** and click **Edit** in the Setup section of the window.
2. Select the **eDirectory** tab and then complete the following fields:
   - **Search Base**—The starting point or root context for agent queries, for example: `dc=domain1, dc=example, dc=com`.
   - **Bind Distinguished Name**—The account to use to bind to the directory, for example: `cn=admin, ou=IT, dc=domain1, dc=example, dc=com`.
   - **Bind Password**—The bind account password. The agent saves the encrypted password in the configuration file.
   - **Search Filter**—The search query for user entries (default is `objectClass=Person`).
   - **Server Domain Prefix**—A prefix to uniquely identify the user. This is only required if there are overlapping name spaces, such as different users with the same name from two different directories.
   - **Use SSL**—Select the check box to use SSL for eDirectory binding.
   - **Verify Server Certificate**—Select the check box to verify the eDirectory server certificate when using SSL.

STEP 4 | (Optional, not recommended) Configure client probing.

Do not enable client probing on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured.

1. On the **Client Probing** tab, select the **Enable WMI Probing** check box and/or the **Enable NetBIOS Probing** check box.
2. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client.

© 2018 Palo Alto Networks, Inc.

For NetBIOS probing to work effectively, each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled. Although client probing is not recommended, if you plan to enable it, WMI probing is preferred over NetBIOS whenever possible.

STEP 5 | Save the configuration.

Click **OK** to save the User-ID agent setup settings and then click **Commit** to restart the User-ID agent and load the new settings.

STEP 6 | (Optional) Define the set of users for which you do not need to provide IP address-to-username mappings, such as kiosk accounts.

You can also use the `ignore-user` list to identify users whom you want to force to authenticate using Captive Portal. Save the `ignore-user` list as a text document using the title `ignore_user_list` and use the .txt file extension to save it to the User-ID Agent folder on the domain server where the agent is installed.

List the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Each user account name must be on a separate line. For example:

```
SPAdmin
SPInstall
TFSReport
```

You can use an asterisk as a wildcard character to match multiple usernames, but only as the last character in the entry. For example, `corpdomain\it-admin*` would match all administrators in the `corpdomain` domain whose usernames start with the string `it#admin`.

STEP 7 | Configure the firewall to connect to the User-ID agent.

The firewall can connect to only one Windows-based User-ID agent that is using the User-ID credential service add-on to detect corporate credential submissions. See Configure Credential Detection with the Windows-based User-ID Agent *for more details on how to use this service for credential phishing prevention.*

Complete the following steps on each firewall you want to connect to the User-ID agent to receive user mappings:

1. Select **Device** > **User Identification** > **User-ID Agents** and click **Add**.
2. Enter a **Name** for the User-ID agent.
3. Enter the IP address of the Windows **Host** on which the User-ID Agent is installed.
4. Enter the **Port** number (1-65535) on which the agent will listen for user mapping requests. This value must match the value configured on the User-ID agent. By default, the port is set to 5007 on the firewall and on newer versions of the User-ID agent. However, some older User-ID agent versions use port 2010 as the default.
5. Make sure that the configuration is **Enabled**, then click **OK**.
6. **Commit** the changes.
7. Verify that the **Connected status** displays as connected (a green light).

STEP 8 | Verify that the User-ID agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.

1. Launch the User-ID agent and select **User Identification**.
2. Verify that the agent status shows **Agent is running**. If the Agent is not running, click **Start**.

© 2018 Palo Alto Networks, Inc.

3. To verify that the User-ID agent can connect to monitored servers, make sure the Status for each Server is **Connected**.
4. To verify that the firewalls can connect to the User-ID agent, make sure the Status for each of the Connected Devices is **Connected**.
5. To verify that the User-ID agent is mapping IP addresses to usernames, select **Monitoring** and make sure that the mapping table is populated. You can also **Search** for specific users, or **Delete** user mappings from the list.

Configure User Mapping Using the PAN-OS Integrated User-ID Agent

The following procedure shows how to configure the PAN-OS integrated User-ID agent on the firewall for IP address-to-username mapping. The integrated User-ID agent performs the same tasks as the Windows-based agent with the exception of NetBIOS client probing (WMI probing is supported).

STEP 1 | Create an Active Directory service account for the User-ID agent to access the services and hosts it will monitor for collecting user mapping information.

Create a Dedicated Service Account for the User-ID Agent on page 459.

STEP 2 | Define the servers that the firewall will monitor to collect user mapping information.

Within the total maximum of 100 monitored servers per firewall, you can define no more than 50 syslog senders for any single virtual system.

To collect all the required mappings, the firewall must connect to all servers that your users log in to so it can monitor the Security log files on all servers that contain login events.

1. Select **Device** > **User Identification** > **User Mapping**.
2. Click **Add** in the Server Monitoring section.
3. Enter a **Name** to identify the server.
4. Select the **Type** of server.
5. Enter the **Network Address** (an FQDN or IP address) of the server.
6. Make sure the server profile is **Enabled** and click **OK**.
7. (Optional) Click **Discover** if you want the firewall to automatically discover domain controllers on your network using DNS lookups.

   The auto-discovery feature is for domain controllers only; you must manually add any Exchange servers or eDirectory servers you want to monitor.
8. (Optional) Specify the frequency at which the firewall polls Windows servers for mapping information. This is the interval between the end of the last query and the start of the next query.

   If the query load is high, the observed delay between queries might significantly exceed the specified frequency.

   1. Edit the Palo Alto Networks User ID Agent Setup.
   2. Select the **Server Monitor** tab and specify the **Server Log Monitor Frequency** in seconds (default is 2, range is 1-3600). Increase the value in this field to 5 seconds in environments with older domain controllers or high-latency links.

      Ensure that the Enable Session setting is not selected. This setting requires that the User-ID agent have an Active Directory account with Server Operator privileges

© 2018 Palo Alto Networks, Inc.

so that it can read all user sessions. Instead, use a Syslog or XML API integration to monitor sources that capture login and logout events for all device types and operating systems (instead of just Windows), such as wireless controllers and NACs.

3. Click **OK** to save the changes.

STEP 3 | Specify the subnetworks the PAN-OS integrated User-ID agent should include in or exclude from user mapping.

By default, the User-ID maps all users accessing the servers you are monitoring.



As a best practice, always specify which networks to include and, optionally, to exclude from User-ID to ensure that the agent is only communicating with internal resources and to prevent unauthorized users from being mapped. You should only enable user mapping on the subnetworks where users internal to your organization are logging in.

1. Select **Device** > **User Identification** > **User Mapping**.
2. **Add** an entry to the Include/Exclude Networks and enter a **Name** for the entry and make sure to keep the **Enabled** check box selected.
3. Enter the **Network Address** and then select whether to include or exclude it:
   - **Include**—Select this option if you want to limit user mapping to users logged in to the specified subnetwork only. For example, if you include 10.0.0.0/8, the agent maps the users on that subnetwork and excludes all others. If you want the agent to map users in other subnetworks, you must repeat these steps to add additional networks to the list.
   - **Exclude**—Select this option only if you want the agent to exclude a subset of the subnetworks you added for inclusion. For example, if you include 10.0.0.0/8 and exclude 10.2.50.0/22, the agent will map users on all the subnetworks of 10.0.0.0/8 except 10.2.50.0/22, and will exclude all subnetworks outside of 10.0.0.0/8.

   If you add subnetworks for exclusion without adding any for inclusion, the agent will not perform user mapping in any subnetwork.
4. Click **OK**.

STEP 4 | Set the domain credentials for the account the firewall will use to access Windows resources. This is required for monitoring Exchange servers and domain controllers as well as for WMI probing.

1. Edit the Palo Alto Networks User ID Agent Setup.
2. Select the **WMI Authentication** tab and enter the **User Name** and **Password** for the account that the User-ID agent will use to probe the clients and monitor servers. Enter the username using the domain \username syntax.

STEP 5 | (Optional, not recommended) Configure WMI probing (the PAN-OS integrated User-ID agent does not support NetBIOS probing).



Do not enable WMI probing on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured.

1. Select the **Client Probing** tab and select the **Enable Probing** check box.
2. (Optional) Modify the **Probe Interval** (in minutes) if necessary to ensure it is long enough for the User-ID agent to probe all the learned IP addresses (default is 20, range is 1-1440). This is the interval between the end of the last probe request and the start of the next request.

© 2018 Palo Alto Networks, Inc.

If the request load is high, the observed delay between requests might significantly exceed the specified interval.

3. Click **OK**.
4. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client.

STEP 6 | (Optional) Define the set of users for which you don't require IP address-to-username mappings, such as kiosk accounts.

You can also use the ignore user list to identify users whom you want to force to authenticate using Captive Portal.

Select the **Ignore User List** tab and **Add** each username to exclude from user mapping. You can use an asterisk as a wildcard character to match multiple usernames, but only as the last character in the entry. For example, **`corpdomain\it-admin*`** would match all administrators in the `corpdomain` domain whose usernames start with the string `it#admin`. You can add up to 5,000 entries to exclude from user mapping.

STEP 7 | Activate your configuration changes.

Click **OK** and **Commit**.

STEP 8 | Verify the configuration.

1. Access the firewall CLI.
2. Enter the following operational command:

```
> show user server-monitor state all
```

3. On the **Device** > **User Identification** > **User Mapping** tab in the web interface, verify that the Status of each server you configured for server monitoring is Connected.

Configure User-ID to Monitor Syslog Senders for User Mapping

To obtain IP address-to-username mappings from existing network services that authenticate users, you can configure the PAN-OS integrated User-ID agent or Windows-based User-ID agent to parse Syslog on page 450 messages from those services. To keep user mappings up to date, you can also configure the User-ID agent to parse syslog messages for logout events so that the firewall automatically deletes outdated mappings.

- Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener on page 474
- Configure the Windows User-ID Agent as a Syslog Listener on page 478

Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener

To configure the PAN-OS Integrated User-ID agent to create new user mappings and remove outdated mappings through syslog monitoring, start by defining Syslog Parse profiles. The User-ID agent uses the profiles to find login and logout events in syslog messages. In environments where *syslog senders* (the network services that authenticate users) deliver syslog messages in different formats, configure a profile for each syslog format. Syslog messages must meet certain criteria for a User-ID agent to parse them (see Syslog). This procedure uses examples with the following formats:

- **Login events**—`[Tue Jul 5 13:15:04 2016 CDT] Administrator authentication success User:johndoe1 Source:192.168.3.212`

© 2018 Palo Alto Networks, Inc.

- **Logout events**—`[Tue Jul 5 13:18:05 2016 CDT] User logout successful User:johndoe1 Source:192.168.3.212`

After configuring the Syslog Parse profiles, you specify syslog senders for the User-ID agent to monitor.

The PAN-OS integrated User-ID agent accepts syslogs over SSL and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog sender. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages. However, if you must use UDP, make sure that the syslog sender and client are both on a dedicated, secure network to prevent untrusted hosts from sending UDP traffic to the firewall.

STEP 1 | Determine whether there is a predefined Syslog Parse profile for your particular syslog senders.

Palo Alto Networks provides several predefined profiles through Application content updates. The predefined profiles are global to the firewall, whereas custom profiles apply to a single virtual system only.

Any new Syslog Parse profiles in a given content release is documented in the corresponding release note along with the specific regex used to define the filter.

1. Install the latest Applications or Applications and Threats update:
    1. Select **Device** > **Dynamic Updates** and **Check Now**.
    2. **Download** and **Install** any new update.
2. Determine which predefined Syslog Parse profiles are available:
    1. Select **Device** > **User Identification** > **User Mapping** and click **Add** in the Server Monitoring section.
    2. Set the **Type** to **Syslog Sender** and click **Add** in the Filter section. If the Syslog Parse profile you need is available, skip the steps for defining custom profiles.

STEP 2 | Define custom Syslog Parse profiles to create and delete user mappings.

Each profile filters syslog messages to identify either login events (to create user mappings) or logout events (to delete mappings), but no single profile can do both.

1. Review the syslog messages that the syslog sender generates to identify the syntax for login and logout events. This enables you to define the matching patterns when creating Syslog Parse profiles.

    While reviewing syslog messages, also determine whether they include the domain name. If they don't, and your user mappings require domain names, enter the Default Domain Name when defining the syslog senders that the User-ID agent monitors (later in this procedure).
2. Select **Device** > **User Identification** > **User Mapping** and edit the Palo Alto Networks User-ID Agent Setup.
3. Select **Syslog Filters** and **Add** a Syslog Parse profile.
4. Enter a name to identify the **Syslog Parse Profile**.
5. Select the **Type** of parsing to find login or logout events in syslog messages:
    - **Regex Identifier**—Regular expressions.
    - **Field Identifier**—Text strings.

© 2018 Palo Alto Networks, Inc.

The following steps describe how to configure these parsing types.

STEP 3 | (**Regex Identifier parsing only**) Define the regex matching patterns.

If the syslog message contains a standalone space or tab as a delimiter, use **`\s`** *for a space and* **`\t`** *for a tab.*

1. Enter the **Event Regex** for the type of events you want to find:
   - **Login events**—For the example message, the regex **`(authentication\ success){1}`** extracts the first **`{1}`** instance of the string `authentication success`.
   - **Logout events**—For the example message, the regex **`(logout\ successful){1}`** extracts the first **`{1}`** instance of the string `logout successful`.

   The backslash (\) before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character.
2. Enter the **Username Regex** to identify the start of the username.

   In the example message, the regex **`User:([a-zA-Z0-9\\\._]+)`** matches the string `User:johndoe1` and identifies `johndoe1` as the username.
3. Enter the **Address Regex** to identify the IP address portion of syslog messages.

   In the example message, the regular expression **`Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})`** matches the IPv4 address `Source:192.168.3.212`.

   The following is an example of a completed Syslog Parse profile that uses regex to identify login events:



4. Click **OK** twice to save the profile.

STEP 4 | (**Field Identifier parsing only**) Define string matching patterns.

1. Enter an **Event String** to identify the type of events you want to find.
   - **Login events**—For the example message, the string `authentication success` identifies login events.
   - **Logout events**—For the example message, the string `logout successful` identifies logout events.
2. Enter a **Username Prefix** to identify the start of the username field in syslog messages. The field does not support regex expressions such as \s (for a space) or \t (for a tab).

   In the example messages, `User:` identifies the start of the username field.
3. Enter the **Username Delimiter** that indicates the end of the username field in syslog messages. Use **`\s`** to indicate a standalone space (as in the sample message) and **`\t`** to indicate a tab.
4. Enter an **Address Prefix** to identify the start of the IP address field in syslog messages. The field does not support regex expressions such as \s (for a space) or \t (for a tab).

   In the example messages, `Source:` identifies the start of the address field.

© 2018 Palo Alto Networks, Inc.

5. Enter the **Address Delimiter** that indicates the end of the IP address field in syslog messages.

   For example, enter **\n** to indicate the delimiter is a line break.

   The following is an example of a completed Syslog Parse profile that uses string matching to identify login events:



6. Click **OK** twice to save the profile.

STEP 5 | Specify the syslog senders that the firewall monitors.

Within the total maximum of 100 monitored servers per firewall, you can define no more than 50 syslog senders for any single virtual system.

The firewall discards any syslog messages received from senders that are not on this list.

1. Select **Device** > **User Identification** > **User Mapping** and **Add** an entry to the Server Monitoring list.
2. Enter a **Name** to identify the sender.
3. Make sure the sender profile is **Enabled** (default is enabled).
4. Set the **Type** to **Syslog Sender**.
5. Enter the **Network Address** of the syslog sender (IP address or FQDN).
6. Select **SSL** (default) or **UDP** as the **Connection Type**.

   Use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog sender. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages when using agentless User Mapping on a firewall. However, if you must use UDP, make sure that the syslog sender and client are both on a dedicated, secure network to prevent untrusted hosts from sending UDP traffic to the firewall.

   A syslog sender using SSL to connect will show a Status of Connected only when there is an active SSL connection. Syslog senders using UDP will not show a Status value.
7. For each syslog format that the sender supports, **Add** a Syslog Parse profile to the Filter list. Select the **Event Type** that each profile is configured to identify: **login** (default) or **logout**.
8. (Optional) If the syslog messages don't contain domain information and your user mappings require domain names, enter a **Default Domain Name** to append to the mappings.
9. Click **OK** to save the settings.

STEP 6 | Enable syslog listener services on the interface that the firewall uses to collect user mappings.

1. Select **Network** > **Network Profiles** > **Interface Mgmt** and edit an existing Interface Management profile or **Add** a new profile.

© 2018 Palo Alto Networks, Inc.

2. Select **User-ID Syslog Listener-SSL** or **User-ID Syslog Listener-UDP** or both, based on the protocols you defined for the syslog senders in the Server Monitoring list.

   The listening ports (514 for UDP and 6514 for SSL) are not configurable; they are enabled through the management service only.
3. Click **OK** to save the interface management profile.

   Even after enabling the User-ID Syslog Listener service on the interface, the interface only accepts syslog connections from senders that have a corresponding entry in the User-ID monitored servers configuration. The firewall discards connections or messages from senders that are not on the list.
4. Assign the Interface Management profile to the interface that the firewall uses to collect user mappings:
   1. Select **Network** > **Interfaces** and edit the interface.
   2. Select **Advanced** > **Other info**, select the Interface **Management Profile** you just added, and click **OK**.
5. **Commit** your changes.

STEP 7 | Verify that the firewall adds and deletes user mappings when users log in and out.

You can use CLI commands to see additional information about syslog senders, syslog messages, and user mappings.

1. Log in to a client system for which a monitored syslog sender generates login and logout event messages.
2. Log in to the firewall CLI.
3. Verify that the firewall mapped the login username to the client IP address:

```
> show user ip-user-mapping ip <ip-address>
IP address:    192.0.2.1 (vsys1)
User:          localdomain\username
From:          SYSLOG
```

4. Log out of the client system.
5. Verify that the firewall deleted the user mapping:

```
> show user ip-user-mapping ip <ip-address>
No matched record
```

Configure the Windows User-ID Agent as a Syslog Listener

To configure the Windows-based User-ID agent to create new user mappings and remove outdated mappings through syslog monitoring, start by defining Syslog Parse profiles. The User-ID agent uses the profiles to find login and logout events in syslog messages. In environments where *syslog senders* (the network services that authenticate users) deliver syslog messages in different formats, configure a profile for each syslog format. Syslog messages must meet certain criteria for a User-ID agent to parse them (see Syslog). This procedure uses examples with the following formats:

- Login events—`[Tue Jul 5 13:15:04 2016 CDT] Administrator authentication success User:johndoe1 Source:192.168.3.212`
- Logout events—`[Tue Jul 5 13:18:05 2016 CDT] User logout successful User:johndoe1 Source:192.168.3.212`

© 2018 Palo Alto Networks, Inc.

After configuring the Syslog Parse profiles, you specify the syslog senders that the User-ID agent monitors.

The Windows User-ID agent accepts syslogs over TCP and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog sender. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, use TCP instead of UDP. In either case, make sure that the syslog sender and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

STEP 1 | Deploy the Windows-based User-ID agents if you haven't already.

1. Install the Windows-Based User-ID Agent.
2. Step Configure the firewall to connect to the User-ID agent.

STEP 2 | Define custom Syslog Parse profiles to create and delete user mappings.

Each profile filters syslog messages to identify either login events (to create user mappings) or logout events (to delete mappings), but no single profile can do both.

1. Review the syslog messages that the syslog sender generates to identify the syntax for login and logout events. This enables you to define the matching patterns when creating Syslog Parse profiles.

   While reviewing syslog messages, also determine whether they include the domain name. If they don't, and your user mappings require domain names, enter the Default Domain Name when defining the syslog senders that the User-ID agent monitors (later in this procedure).
2. Open the Windows **Start** menu and select **User-ID Agent**.
3. Select **User Identification** > **Setup** and **Edit** the Setup.
4. Select **Syslog**, **Enable Syslog Service**, and **Add** a Syslog Parse profile.
5. Enter a **Profile Name** and **Description**.
6. Select the **Type** of parsing to find login and logout events in syslog messages:
   - **Regex**—Regular expressions.
   - **Field**—Text strings.

   The following steps describe how to configure these parsing types.

STEP 3 | (**Regex parsing only**) Define the regex matching patterns.

If the syslog message contains a standalone space or tab as a delimiter, use **`\s`** for a space and **`\t`** for a tab.

1. Enter the **Event Regex** for the type of events you want to find:
   - **Login events**—For the example message, the regex **`(authentication\ success){1}`** extracts the first **`{1}`** instance of the string `authentication success`.
   - **Logout events**—For the example message, the regex **`(logout\ successful){1}`** extracts the first **`{1}`** instance of the string `logout successful`.

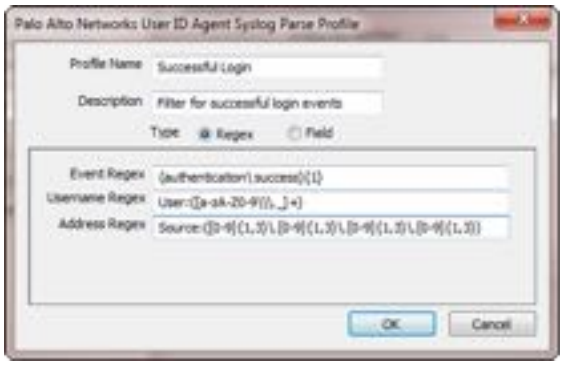
   The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character.
2. Enter the **Username Regex** to identify the start of the username.

   In the example message, the regex **`User:([a-zA-Z0-9\\\._]+)`** matches the string `User:johndoe1` and identifies `johndoe1` as the username.
3. Enter the **Address Regex** to identify the IP address portion of syslog messages.

© 2018 Palo Alto Networks, Inc.

In the example message, the regular expression **`Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})`** matches the IPv4 address `Source:192.168.3.212`.

The following is an example of a completed Syslog Parse profile that uses regex to identify login events:



4. Click **OK** twice to save the profile.

STEP 4 | (**Field Identifier parsing only**) Define string matching patterns.

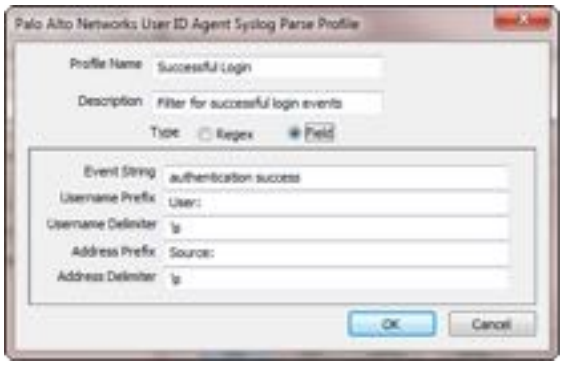
1. Enter an **Event String** to identify the type of events you want to find.
   - **Login events**—For the example message, the string `authentication success` identifies login events.
   - **Logout events**—For the example message, the string `logout successful` identifies logout events.
2. Enter a **Username Prefix** to identify the start of the username field in syslog messages. The field does not support regex expressions such as \s (for a space) or \t (for a tab).

   In the example messages, `User:` identifies the start of the username field.
3. Enter the **Username Delimiter** that indicates the end of the username field in syslog messages. Use **`\s`** to indicate a standalone space (as in the sample message) and **`\t`** to indicate a tab.
4. Enter an **Address Prefix** to identify the start of the IP address field in syslog messages. The field does not support regex expressions such as \s (for a space) or \t (for a tab).

   In the example messages, `Source:` identifies the start of the address field.
5. Enter the **Address Delimiter** that indicates the end of the IP address field in syslog messages.

   For example, enter **`\n`** to indicate the delimiter is a line break.

   The following is an example of a completed Syslog Parse profile that uses string matching to identify login events:



6. Click **OK** twice to save the profile.

© 2018 Palo Alto Networks, Inc.

STEP 5 | Specify the syslog senders that the User-ID agent monitors.

Within the total maximum of 100 servers of all types that the User-ID agent can monitor, up to 50 can be syslog senders.

The User-ID agent discards any syslog messages received from senders that are not on this list.

1. Select **User Identification** > **Discovery** and **Add** an entry to the Servers list.
2. Enter a **Name** to identify the sender.
3. Enter the **Server Address** of the syslog sender (IP address or FQDN).
4. Set the **Server Type** to **Syslog Sender**.
5. (**Optional**) If you want to override the current domain in the username of your syslog message or prepend the domain to the username if your syslog message doesn't contain a domain, enter a **Default Domain Name**.
6. For each syslog format that the sender supports, **Add** a Syslog Parse profile to the Filter list. Select the **Event Type** that you configured each profile to identify—**login** (default) or **logout**—and then click **OK**.
7. Click **OK** to save the settings.
8. **Commit** your changes to the User-ID agent configuration.

STEP 6 | Verify that the User-ID agent adds and deletes user mappings when users log in and out.



You can use CLI commands to see additional information about syslog senders, syslog messages, and user mappings.

1. Log in to a client system for which a monitored syslog sender generates login and logout event messages.
2. Verify that the User-ID agent mapped the login username to the client IP address:
   1. In the User-ID agent, select **Monitoring**.
   2. Enter the username or IP address in the filter field, **Search**, and verify that the list displays the mapping.
3. Verify that the firewall received the user mapping from the User-ID agent:
   1. Log in to the firewall CLI.
   2. Run the following command:

      ```
      > show user ip-user-mapping ip <ip-address>
      ```

      If the firewall received the user mapping, the output resembles the following:

      ```
      IP address:    192.0.2.1 (vsys1)
      User:          localdomain\username
      From:          SYSLOG
      ```
4. Log out of the client system.
5. Verify that the User-ID agent removed the user mapping:
   1. In the User-ID agent, select **Monitoring**.
   2. Enter the username or IP address in the filter field, **Search**, and verify that the list does not display the mapping.
6. Verify that the firewall deleted the user mapping:
   1. Access the firewall CLI.
   2. Run the following command:

© 2018 Palo Alto Networks, Inc.

```
> show user ip-user-mapping ip <ip-address>
```

If the firewall deleted the user mapping, the output displays:

```
No matched record
```

Map IP Addresses to Usernames Using Captive Portal

When a user initiates web traffic (HTTP or HTTPS) that matches an Authentication Policy on page 211 rule, the firewall prompts the user to authenticate through Captive Portal. This ensures that you know exactly who is accessing your most sensitive applications and data. Based on user information collected during authentication, the firewall creates a new IP address-to-username mapping or updates the existing mapping for that user. This method of user mapping is useful in environments where the firewall cannot learn mappings through other methods such as monitoring servers. For example, you might have users who are not logged in to your monitored domain servers, such as users on Linux clients.

- Captive Portal Authentication Methods on page 482
- Captive Portal Modes on page 483
- Configure Captive Portal on page 484

Captive Portal Authentication Methods

Captive Portal uses the following methods to authenticate users whose web requests match Authentication Policy rules:

Authentication Method	Description
Kerberos SSO	The firewall uses Kerberos single sign-on (SSO) to transparently obtain user credentials from the browser. To use this method, your network requires a Kerberos infrastructure, including a key distribution center (KDC) with an authentication server and ticket granting service. The firewall must have a Kerberos account. If Kerberos SSO authentication fails, the firewall falls back to NT LAN Manager (NTLM) authentication. If you don't configure NTLM, or NTLM authentication fails, the firewall falls back to web form or client certificate authentication, depending on your Authentication policy and Captive Portal configuration.  *Kerberos SSO is preferable to NTLM authentication. Kerberos is a stronger, more robust authentication method than NTLM and it does not require the firewall to have an administrative account to join the domain.*
NT LAN Manager (NTLM)	The firewall uses an encrypted challenge-response mechanism to obtain the user credentials from the browser. When configured properly, the browser will transparently provide the credentials to the firewall without prompting the user, but will prompt for credentials if necessary. If you use the Windows-based User-ID agent, NTLM responses go directly to the domain controller where you installed the agent.

© 2018 Palo Alto Networks, Inc.

Authentication Method	Description
	If you configure Kerberos SSO authentication, the firewall tries that method first before falling back to NTLM authentication. If the browser can't perform NTLM or if NTLM authentication fails, the firewall falls back to web form or client certificate authentication, depending on your Authentication policy and Captive Portal configuration. Microsoft Internet Explorer supports NTLM by default. You can configure Mozilla Firefox and Google Chrome to also use NTLM but you can't use NTLM to authenticate non-Windows clients.
Web Form	The firewall redirects web requests to a web form for authentication. For this method, you can configure Authentication policy to use Multi-Factor Authentication (MFA), SAML, Kerberos, TACACS+, RADIUS, or LDAP authentication. Although users have to manually enter their login credentials, this method works with all browsers and operating systems.
Client Certificate Authentication	The firewall prompts the browser to present a valid client certificate to authenticate the user. To use this method, you must provision client certificates on each user system and install the trusted certificate authority (CA) certificate used to issue those certificates on the firewall.

Captive Portal Modes

The Captive Portal mode defines how the firewall captures web requests for authentication:

Mode	Description
Transparent	The firewall intercepts the browser traffic per the Authentication policy rule and impersonates the original destination URL, issuing an HTTP 401 to invoke authentication. However, because the firewall does not have the real certificate for the destination URL, the browser displays a certificate error to users attempting to access a secure site. Therefore, use this mode only when absolutely necessary, such as in Layer 2 or virtual wire deployments.
Redirect	The firewall intercepts unknown HTTP or HTTPS sessions and redirects them to a Layer 3 interface on the firewall using an HTTP 302 redirect to perform authentication. This is the preferred mode because it provides a better end-user experience (no certificate errors). However, it does require additional Layer 3 configuration. Another benefit of the Redirect mode is that it provides for the use of session cookies, which enable the user to continue browsing to authenticated sites without requiring re-mapping each time the timeouts expire. This is especially useful for users who roam from one IP address to another (for example, from the corporate LAN to the wireless network) because they won't need to re-authenticate when the IP address changes as long as the session stays open. If you use Kerberos SSO or NTLM authentication, you must use Redirect mode because the browser will provide credentials only to

© 2018 Palo Alto Networks, Inc.

Mode	Description
	trusted sites. Redirect mode is also required if you use Multi-Factor Authentication on page 155 to authenticate Captive Portal users.

Configure Captive Portal

The following procedure shows how to set up Captive Portal authentication by configuring the PAN-OS integrated User-ID agent to redirect web requests that match an Authentication Policy rule to a firewall interface (redirect host). Based on their sensitivity, the applications that users access through Captive Portal require different authentication methods and settings. To accommodate all authentication requirements, you can use default and custom authentication enforcement objects. Each object associates an Authentication rule with an authentication profile and a Captive Portal authentication method.

- **Default authentication enforcement objects**—Use the default objects if you want to associate multiple Authentication rules with the same global authentication profile. You must configure this authentication profile before configuring Captive Portal, and then assign it in the Captive Portal Settings. For Authentication rules that require Multi-Factor Authentication (MFA), you cannot use default authentication enforcement objects.
- **Custom authentication enforcement objects**—Use a custom object for each Authentication rule that requires an authentication profile that differs from the global profile. Custom objects are mandatory for Authentication rules that require MFA. To use custom objects, create authentication profiles and assign them to the objects after configuring Captive Portal—when you Configure Authentication Policy.

Keep in mind that authentication profiles are necessary only if users authenticate through a Captive Portal Web Form, Kerberos SSO, or NT LAN Manager (NTLM). Alternatively, or in addition to these methods, the following procedure also describes how to implement Client Certificate Authentication.

If you use Captive Portal without the other User-ID functions (user mapping and group mapping), you don't need to configure a User-ID agent.

STEP 1 | Configure the interfaces that the firewall will use for incoming web requests, authenticating users, and communicating with directory servers to map usernames to IP addresses.

By default, the firewall uses the management interface for authentication requests and communication with the Active Directory and User-ID agent. You can redirect Captive Portal requests to a dataplane interface by configuring service routes.

1. (MGT interface only) Select **Device** > **Setup** > **Interfaces**, edit the **Management** interface, select **User-ID**, and click **OK**.
2. (Non-MGT interface only) Assign an Interface Management Profile to the Layer 3 interface that the firewall will use for incoming web requests and communication with directory servers. You must enable **Response Pages** and **User-ID** in the Interface Management profile.
3. (Non-MGT interface only) Configure a service route for the interface that the firewall will use to authenticate users. If the firewall has more than one virtual system (vsys), the service route can be global or vsys-specific. The services must include **LDAP** and potentially the following:
   - **Kerberos**, **RADIUS**, **TACACS+**, or **Multi-Factor Authentication**—Configure a service route for any authentication services that you use.
   - **UID Agent**—Configure this service only if you will enable NT LAN Manager (NTLM) authentication or if you will Enable User- and Group-Based Policy.
4. (Redirect mode only) Create a DNS address (A) record that maps the IP address on the Layer 3 interface to the redirect host. If you will use Kerberos SSO, you must also add a DNS pointer (PTR) record that performs the same mapping.

© 2018 Palo Alto Networks, Inc.

If your network doesn't support access to the directory servers from any firewall interface, you must Configure User Mapping Using the Windows User-ID Agent.

STEP 2 | Make sure Domain Name System (DNS) is configured to resolve your domain controller addresses.

To verify proper resolution, ping the server FQDN. For example:

```
admin@PA-200> ping host dc1.acme.com
```

STEP 3 | Configure clients to trust Captive Portal certificates.

Required for redirect mode—to transparently redirect users without displaying certificate errors. You can generate a self-signed certificate or import a certificate that an external certificate authority (CA) signed.

To use a self-signed certificate, create a root CA certificate and use it to sign the certificate you will use for Captive Portal:

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
2. Create a Self-Signed Root CA Certificate or import a CA certificate (see Import a Certificate and Private Key).
3. Generate a Certificate to use for Captive Portal. Be sure to configure the following fields:
   - **Common Name**—Enter the DNS name of the intranet host for the Layer 3 interface.
   - **Signed By**—Select the CA certificate you just created or imported.
   - Certificate Attributes—Click **Add**, for the **Type** select **IP** and, for the **Value**, enter the IP address of the Layer 3 interface to which the firewall will redirect requests.
4. Configure an SSL/TLS Service Profile. Assign the Captive Portal certificate you just created to the profile.
5. Configure clients to trust the certificate:
   1. Export the CA certificate you created or imported.
   2. Import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory (AD) Group Policy Object (GPO).

STEP 4 | (Optional) Configure Client Certificate Authentication.



You don't need an authentication profile or sequence for client certificate authentication. If you configure both an authentication profile/sequence and certificate authentication, users must authenticate using both.

1. Use a root CA certificate to generate a client certificate for each user who will authenticate through Captive Portal. The CA in this case is usually your enterprise CA, not the firewall.
2. Export the CA certificate in PEM format to a system that the firewall can access.
3. Import the CA certificate onto the firewall: see Import a Certificate and Private Key. After the import, click the imported certificate, select **Trusted Root CA**, and click **OK**.
4. Configure a Certificate Profile.
   - In the **Username Field** drop-down, select the certificate field that contains the user identity information.
   - In the **CA Certificates** list, click **Add** and select the CA certificate you just imported.

STEP 5 | (Optional) Enable NT LAN Manager (NTLM) authentication.

© 2018 Palo Alto Networks, Inc.



As a best practice, choose Kerberos single sign-on (SSO) or SAML SSO authentication over NTLM authentication. Kerberos and SAML are stronger, more robust authentication methods than NTLM and do not require the firewall to have an administrative account to join the domain. If you do configure NTLM, the PAN-OS integrated User-ID agent must be able to successfully resolve the DNS name of your domain controller to join the domain.

1. If you haven't already done so, Create a Dedicated Service Account for the User-ID Agent.

   As a best practice, you use a User-ID agent account that is separate from your firewall administrator account.
2. Select **Device** > **User Identification** > **User Mapping** and edit the Palo Alto Networks User ID Agent Setup section.
3. Select **NTLM** and **Enable NTLM authentication processing**.
4. Enter the **NTLM Domain** against which the User-ID agent on the firewall will check NTLM credentials.
5. Enter the **Admin User Name** and **Password** of the Active Directory account you created for the User-ID agent.

   Do not include the domain in the Admin User Name field. Otherwise, the firewall will fail to join the domain.
6. Click **OK** to save your settings.

STEP 6 | (Optional) Configure Captive Portal for the Apple Captive Network Assistant.

This step is only required if you are using Captive Portal with the Apple Captive Network Assistant (CNA). To use Captive Portal with CNA, perform the following steps.

1. Verify you have specified an FQDN for the redirect host (not just an IP address).
2. Select an SSL/TLS service profile that uses a publicly-signed certificate for the specified FQDN.
3. Enter the following command to adjust the number of requests supported for Captive Portal: **`set deviceconfig setting ctd cap-portal-ask-requests <threshold-value>`**

   By default, the firewall has a rate limit threshold for Captive Portal that limits the number of requests to one request every two seconds. The CNA sends multiple requests that can exceed this limit, which can result in a TCP reset and an error from the CNA. The recommended threshold value is 5 (default is one). This value will allow up to 5 requests every two seconds. Based on your environment, you may need to configure a different value. If the current value is not sufficient to handle the number of requests, increase the value.

STEP 7 | Configure the Captive Portal settings.

1. Select **Device** > **User Identification** > **Captive Portal Settings** and edit the settings.
2. **Enable Captive Portal** (default is enabled).
3. Specify the **Timer**, which is the maximum time in minutes that the firewall retains an IP address-to-username mapping for a user after that user authenticates through Captive Portal (default is 60; range is 1 to 1,440). After the **Timer** expires, the firewall removes the mapping and any associated Authentication Timestamps used to evaluate the **Timeout** in Authentication policy rules.

   When evaluating the Captive Portal Timer and the Timeout value in each Authentication policy rule, the firewall prompts the user to re-authenticate for whichever setting expires first. Upon re-authenticating, the firewall resets the time count for the Captive Portal Timer and records new authentication timestamps for the user. Therefore, to enable different Timeout periods for different Authentication rules, set the Captive Portal Timer to a value the same as or higher than any rule Timeout.

© 2018 Palo Alto Networks, Inc.

4. Select the **SSL/TLS Service Profile** you created for redirect requests over TLS. See Configure an SSL/TLS Service Profile.
5. Select the **Mode** (in this example, **Redirect**).
6. (Redirect mode only) Specify the **Redirect Host**, which is the intranet hostname (a hostname with no period in its name) that resolves to the IP address of the Layer 3 interface on the firewall to which web requests are redirected.

   If users authenticate through Kerberos single sign-on (SSO), the **Redirect Host** must be the same as the hostname specified in the Kerberos keytab.
7. Select the authentication method to use if NTLM fails (or if you don't use NTLM):
   - To use client certificate authentication, select the **Certificate Profile** you created.
   - To use global settings for interactive or SSO authentication, select the **Authentication Profile** you configured.
   - To use Authentication policy rule-specific settings for interactive or SSO authentication, assign authentication profiles to authentication enforcement objects when you Configure Authentication Policy.
8. Click **OK** and **Commit** the Captive Portal configuration.

STEP 8 | Next steps...

The firewall does not display the Captive Portal web form to users until you Configure Authentication Policy rules that trigger authentication when users request services or applications.

Configure User Mapping for Terminal Server Users

Individual terminal server users appear to have the same IP address and therefore an IP address-to-username mapping is not sufficient to identify a specific user. To enable identification of specific users on Windows-based terminal servers, the Palo Alto Networks Terminal Services agent (TS agent) allocates a port range to each user. It then notifies every connected firewall about the allocated port range, which allows the firewall to create an IP address-port-user mapping table and enable user- and group-based security policy enforcement. For non-Windows terminal servers, you can configure the PAN-OS XML API to extract user mapping information.

The following sections describe how to configure user mapping for terminal server users:

- Configure the Palo Alto Networks Terminal Services Agent for User Mapping on page 487
- Retrieve User Mappings from a Terminal Server Using the PAN-OS XML API on page 491

Configure the Palo Alto Networks Terminal Services Agent for User Mapping

Use the following procedure to install and configure the TS agent on the terminal server. To map all your users, you must install the TS agent on all terminal servers that your users log in to.



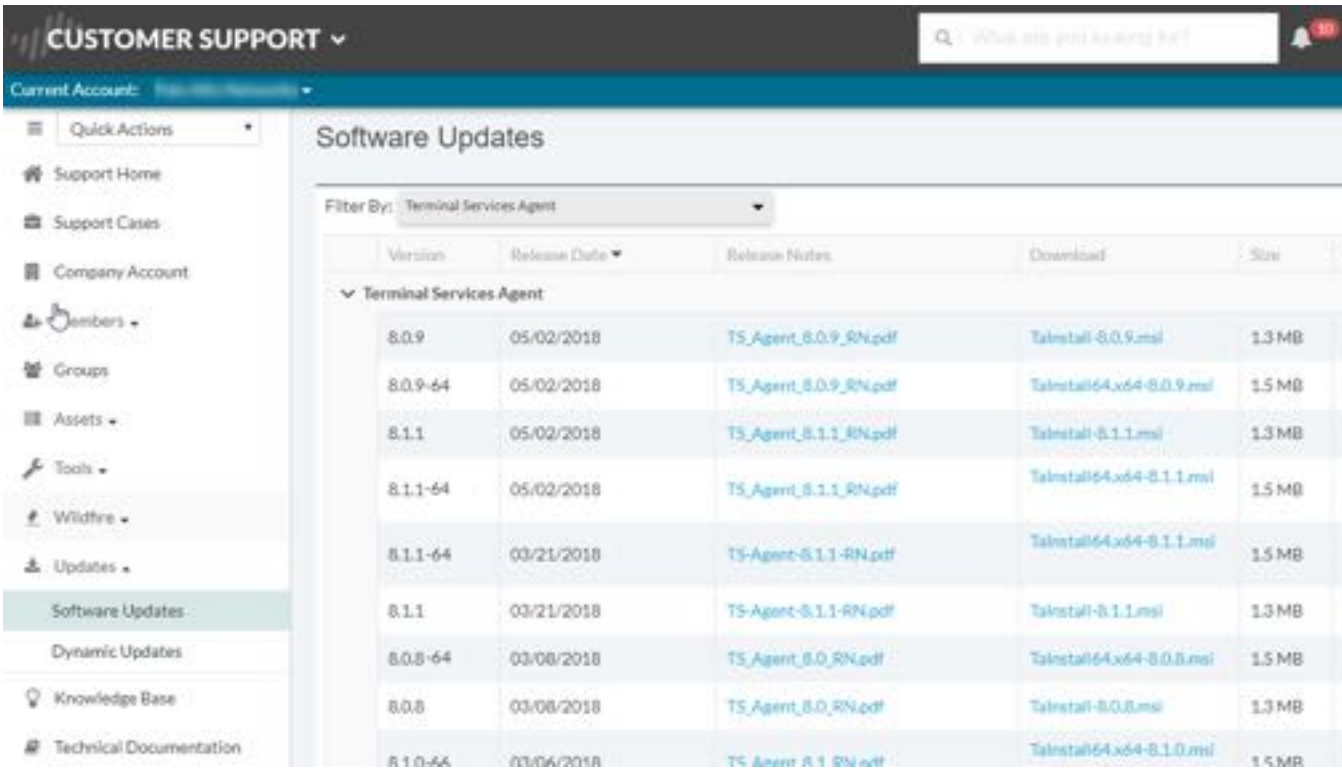
- *If you are using version 7.0 or later of the TS agent, make sure to disable any Sophos anti-virus software on the TS agent host. Otherwise, the source ports allocated by the TS agent will be overwritten.*
- *If you are installing the TS agent version 8.0 or later on a Windows Server 2008 R2, download and install the security advisory patch on the server before installing the TS agent. Otherwise, the TS agent service fails to launch with the following error written to the logs:* `service fails with error 577`.

For information about the supported terminal servers supported by the TS Agent, refer to the Palo Alto Networks Compatibility Matrix.

© 2018 Palo Alto Networks, Inc.

STEP 1 | Download the TS agent installer.

1. Log in to the Palo Alto Networks Customer Support Portal.
2. Select **Updates** > **Software Updates**.
3. Set **Filter By** to **Terminal Services Agent** and select the version of the agent you want to install from the corresponding Download column. For example, to download the 8.0 version of the agent, select **TaInstall-8.0.0.msi**.
4. Save the `TaInstall.x64-x.x.x-xx.msi` or `TaInstall-x.x.x-xx.msi` file (be sure to select the appropriate version based on whether the Windows system is running a 32-bit OS or a 64-bit OS) on the systems where you plan to install the agent.



STEP 2 | Run the installer as an administrator.

1. Open the Windows **Start** menu, right-click the **Command Prompt** program, and select **Run as administrator**.
2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop you would enter the following:

```
C:\Users\administrator.acme>cd Desktop
C:\Users\administrator.acme\Desktop>TaInstall-8.0.0-1.msi
```

3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to the `C:\Program Files (x86)\Palo Alto Networks\Terminal Server Agent` folder, but you can **Browse** to a different location.
4. When the installation completes, **Close** the setup window.

If you are upgrading to a TS Agent version that has a newer driver than the existing installation, the installation wizard prompts you to reboot the system after upgrading in order to use the new driver.

STEP 3 | Define the range of ports for the TS Agent to allocate to end users.

© 2018 Palo Alto Networks, Inc.

The System Source Port Allocation Range and System Reserved Source Ports fields specify the range of ports that will be allocated to non-user sessions. Make sure the values specified in these fields do not overlap with the ports you designate for user traffic. These values can only be changed by editing the corresponding Windows registry settings.The TS agent does not allocate ports for network traffic emitted by session 0.

1. Open the Windows **Start** menu and select **Terminal Server Agent** to launch the Terminal Services agent application.
2. Select **Configure** in the side menu.
3. Enter the **Source Port Allocation Range** (default 20000-39999). This is the full range of port numbers that the TS agent will allocate for user mapping. The port range you specify cannot overlap with the **System Source Port Allocation Range**.
4. (Optional) If there are ports/port ranges within the source port allocation that you do not want the TS Agent to allocate to user sessions, specify them as **Reserved Source Ports**. To include multiple ranges, use commas with no spaces, for example: **2000–3000,3500,4000–5000**.
5. Specify the number of ports to allocate to each individual user upon login to the terminal server in the **Port Allocation Start Size Per User** field (default 200).
6. Specify the **Port Allocation Maximum Size Per User**, which is the maximum number of ports the Terminal Services agent can allocate to an individual user.
7. Specify whether to continue processing traffic from the user if the user runs out of allocated ports. By default, the **Fail port binding when available ports are used up** is selected, which indicates that the application will fail to send traffic when all ports are used. To enable users to continue using applications when they run out of ports, clear this check box. Keep in mind that this traffic may not be identified with User-ID.

STEP 4 | (Optional) Assign your own certificates for mutual authentication between the TS agent and the firewall.

1. Obtain your certificate for the TS agent for your enterprise PKI or generate one on your firewall. The private key of the server certificate must be encrypted. The certificate must be uploaded in PEM file format.
   - Generate a Certificate and export it for upload to the TS agent.
   - Export a certificate from your enterprise certificate authority (CA) and the upload it to the TS agent.
2. Add a server certificate to TS agent.
   1. On the TS agent, select **Server Certificate** and click **Add**.
   2. Enter the path and name of the certificate file received from the CA or browse to the certificate file.
   3. Enter the private key password.
   4. Click **OK** and then **Commit**.
3. Configure and assign the certificate profile for the firewall.
   1. Select **Device** > **Certificate Management** > **Certificate Profile** to Configure a Certificate Profile.

      You can only assign one certificate profile for Windows User-ID agents and TS agents. Therefore, your certificate profile must include all certificate authorities that issued certificates uploaded to connected Windows User-ID and TS agents.
   2. Select **Device** > **User Identification** > **Connection Security** and click the edit button to assign the certificate profile.
   3. Select the certificate profile you configured in the previous step from the User-ID Certificate Profile drop-down.
   4. Click **OK**.

© 2018 Palo Alto Networks, Inc.

5. **Commit** your changes.

STEP 5 | Configure the firewall to connect to the Terminal Services agent.

Complete the following steps on each firewall you want to connect to the Terminal Services agent to receive user mappings:

1. Select **Device** > **User Identification** > **Terminal Server Agents** and click **Add**.
2. Enter a **Name** for the Terminal Services agent.
3. Enter the IP address of the Windows **Host** on which the Terminal Services agent is installed.



When you configure a TS agent connection, the Best Practice is to set the Host and any Alternative IP Addresses to a static IP address that never changes or to an FQDN that resolves to a static IP address.

4. Enter the **Port** number on which the agent will listen for user mapping requests. This value must match the value configured on the Terminal Services agent. By default, the port is set to 5009 on the firewall and on the agent. If you change it here, you must also change the **Listening Port** field on the Terminal Services agent **Configure** screen.
5. Make sure that the configuration is **Enabled** and then click **OK**.
6. **Commit** the changes.
7. Verify that the **Connected** status displays as connected (a green light).

STEP 6 | Verify that the Terminal Services agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.

1. Open the Windows **Start** menu and select **Terminal Server Agent**.
2. Verify that the firewalls can connect by making sure the **Connection Status** of each firewall in the Connection List is **Connected**.
3. Verify that the Terminal Services agent is successfully mapping port ranges to usernames by selecting **Monitor** in the side menu and making sure that the mapping table is populated.

STEP 7 | (Windows 2012 R2 servers only) Disable Enhanced Protected Mode in Microsoft Internet Explorer for each user who uses that browser.

This task is not necessary for other browsers such as Google Chrome or Mozilla Firefox.



To disable Enhanced Protected Mode for all users, use Local Security Policy*.*

Perform these steps on the Windows Server:

1. Start Internet Explorer.
2. Select **Internet options** > **Advanced** and scroll down to the Security section.
3. Clear **Enable Enhanced Protected Mode**.
4. Click **OK**.



In Internet Explorer, Palo Alto Networks recommends that you do not disable Protected Mode, which differs from Enhanced Protected Mode.

© 2018 Palo Alto Networks, Inc.

Retrieve User Mappings from a Terminal Server Using the PAN-OS XML API

The PAN-OS XML API uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services.

To enable a non-Windows terminal server to send user mapping information directly to the firewall, create scripts that extract the user login and logout events and use them for input to the PAN-OS XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget and providing the firewall's API key for secure communication. Creating user mappings from multi-user systems such as terminal servers requires use of the following API messages:

- **<multiusersystem>**—Sets up the configuration for an XML API Multi-user System on the firewall. This message allows for definition of the terminal server IP address (this will be the source address for all users on that terminal server). In addition, the **<multiusersystem>** setup message specifies the range of source port numbers to allocate for user mapping and the number of ports to allocate to each individual user upon login (called the *block size*). If you want to use the default source port allocation range (1025-65534) and block size (200), you do not need to send a **<multiusersystem>** setup event to the firewall. Instead, the firewall will automatically generate the XML API Multi-user System configuration with the default settings upon receipt of the first user login event message.
- **<blockstart>**—Used with the **<login>** and **<logout>** messages to indicate the starting source port number allocated to the user. The firewall then uses the block size to determine the actual range of port numbers to map to the IP address and username in the login message. For example, if the **<blockstart>** value is 13200 and the block size configured for the multi-user system is 300, the actual source port range allocated to the user is 13200 through 13499. Each connection initiated by the user should use a unique source port number within the allocated range, enabling the firewall to identify the user based on its IP address-port-user mappings for enforcement of user- and group-based security rules. When a user exhausts all the ports allocated, the terminal server must send a new **<login>** message allocating a new port range for the user so that the firewall can update the IP address-port-user mapping. In addition, a single username can have multiple blocks of ports mapped simultaneously. When the firewall receives a **<logout>** message that includes a **<blockstart>** parameter, it removes the corresponding IP address-port-user mapping from its mapping table. When the firewall receives a **<logout>** message with a username and IP address, but no **<blockstart>**, it removes the user from its table. And, if the firewall receives a **<logout>** message with an IP address only, it removes the multi-user system and all mappings associated with it.



The XML files that the terminal server sends to the firewall can contain multiple message types and the messages do not need to be in any particular order within the file. However, upon receiving an XML file that contains multiple message types, the firewall will process them in the following order: multiusersystem requests first, followed by logins, then logouts.

The following workflow provides an example of how to use the PAN-OS XML API to send user mappings from a non-Windows terminal server to the firewall.

STEP 1 | Generate the API key that will be used to authenticate the API communication between the firewall and the terminal server. To generate the key you must provide login credentials for an administrative account; the API is available to all administrators (including role-based administrators with XML API privileges enabled).



Any special characters in the password must be URL/ percent-encoded.

© 2018 Palo Alto Networks, Inc.

From a browser, log in to the firewall. Then, to generate the API key for the firewall, open a new browser window and enter the following URL:

```
https://<Firewall-IPaddress>/api/?
type=keygen&user=<username>&password=<password>
```

Where **<Firewall-IPaddress>** is the IP address or FQDN of the firewall and **<username>** and **<password>** are the credentials for the administrative user account on the firewall. For example:

```
https://10.1.2.5/api/?type=keygen&user=admin&password=admin
```

The firewall responds with a message containing the key, for example:

```
<response status="success">
   <result>
      <key>k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg=</key>
   </result>
</response>
```

STEP 2 | (Optional) Generate a setup message that the terminal server will send to specify the port range and block size of ports per user that your terminal services agent uses.

If the terminal services agent does not send a setup message, the firewall will automatically create a Terminal Services agent configuration using the following default settings upon receipt of the first login message:

- Default port range: 1025 to 65534
- Per user block size: 200
- Maximum number of multi-user systems: 1,000

The following shows a sample setup message:

```
<uid-message>
  <payload>
<multiusersystem>
<entry ip="10.1.1.23" startport="20000"          endport="39999"
 blocksize="100/">
</multiusersystem>
</payload>
<type>update</type>
<version>1.0</version>
</uid-message>
```

where `entry ip` specifies the IP address assigned to terminal server users, `startport` and `endport` specify the port range to use when assigning ports to individual users, and `blocksize` specifies the number of ports to assign to each user. The maximum blocksize is 4000 and each multi-user system can allocate a maximum of 1000 blocks.

If you define a custom blocksize and or port range, keep in mind that you must configure the values such that every port in the range gets allocated and that there are no gaps or unused ports. For example, if you set the port range to 1000–1499, you could set the block size to 100, but not to 200. This is because if you set it to 200, there would be unused ports at the end of the range.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Create a script that will extract the login events and create the XML input file to send to the firewall.

Make sure the script enforces assignment of port number ranges at fixed boundaries with no port overlaps. For example, if the port range is 1000–1999 and the block size is 200, acceptable blockstart values would be 1000, 1200, 1400, 1600, or 1800. Blockstart values of 1001, 1300, or 1850 would be unacceptable because some of the port numbers in the range would be left unused.

The login event payload that the terminal server sends to the firewall can contain multiple login events.

The following shows the input file format for a PAN-OS XML login event:

```
<uid-message>
<payload>
<login>
<entry name="acme\jjaso" ip="10.1.1.23" blockstart="20000">
<entry name="acme\jparker" ip="10.1.1.23" blockstart="20100">
<entry name="acme\ccrisp" ip="10.1.1.23" blockstart="21000">
</login>
</payload>
<type>update</type>
<version>1.0</version>
</uid-message>
```

The firewall uses this information to populate its user mapping table. Based on the mappings extracted from the example above, if the firewall received a packet with a source address and port of 10.1.1.23:20101, it would map the request to user jparker for policy enforcement.

Each multi-user system can allocate a maximum of 1,000 port blocks.

STEP 4 | Create a script that will extract the logout events and create the XML input file to send to the firewall.

Upon receipt of a `logout` event message with a `blockstart` parameter, the firewall removes the corresponding IP address-port-user mapping. If the `logout` message contains a username and IP address, but no `blockstart` parameter, the firewall removes all mappings for the user. If the `logout` message contains an IP address only, the firewall removes the multi-user system and all associated mappings.

The following shows the input file format for a PAN-OS XML logout event:

```
<uid-message>
<payload>
<logout>
<entry name="acme\jjaso" ip="10.1.1.23" blockstart="20000">
<entry name="acme\ccrisp" ip="10.1.1.23">
<entry ip="10.2.5.4">
</logout>
</payload>
<type>update</type>
<version>1.0</version>
</uid-message>
```

You can also clear the multiuser system entry from the firewall using the following CLI command: ***clear xml-api multiusersystem***

© 2018 Palo Alto Networks, Inc.

STEP 5 | Make sure that the scripts you create include a way to dynamically enforce that the port block range allocated using the XML API matches the actual source port assigned to the user on the terminal server and that the mapping is removed when the user logs out or the port allocation changes.

One way to do this would be to use netfilter NAT rules to hide user sessions behind the specific port ranges allocated via the XML API based on the uid. For example, to ensure that a user with the user ID jjaso is mapped to a source network address translation (SNAT) value of 10.1.1.23:20000-20099, the script you create should include the following:

```
[root@ts1 ~]# iptables -t nat -A POSTROUTING -m owner --uid-owner jjaso -p
 tcp -j SNAT --to-source 10.1.1.23:20000-20099
```

Similarly, the scripts you create should also ensure that the IP table routing configuration dynamically removes the SNAT mapping when the user logs out or the port allocation changes:

```
[root@ts1 ~]# iptables -t nat -D POSTROUTING 1
```

STEP 6 | Define how to package the XML input files containing the setup, login, and logout events into wget or cURL messages for transmission to the firewall.

To apply the files to the firewall using wget:

```
> wget --post file <filename> "https://<Firewall-
IPaddress>/api/?type=user-id&key=<key>&file-
name=<input_filename.xml>&client=wget&vsys=<VSYS_name>"
```

For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key `k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg` using wget would look as follows:

```
> wget --post file login.xml "https://10.2.5.11/api/?type=user-
id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg&file-
name=login.xml&client=wget&vsys=vsys1"
```

To apply the file to the firewall using cURL:

```
> curl --form file=@<filename> https://<Firewall-IPaddress>/api/?
type=user-id&key=<key>&vsys=<VSYS_name>
```

For example, the syntax for sending an input file named login.xml to the firewall at 10.2.5.11 using key `k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg` using cURL would look as follows:

```
> curl --form file@login.xml "https://10.2.5.11/api/?type=user-
id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOnlZRg&vsys=vsys1"
```

STEP 7 | Verify that the firewall is successfully receiving login events from the terminal servers.

Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To verify if the terminal server is connecting to the firewall over XML:

```
admin@PA-5050> show user xml-api multiusersystem
Host            Vsys    Users   Blocks
--------------------------------------
10.5.204.43   vsys1      5         2
```


© 2018 Palo Alto Networks, Inc.

To verify that the firewall is receiving mappings from a terminal server over XML:

```
admin@PA-5050> show user ip-port-user-mapping all

Global max host index 1, host hash count 1

XML API Multi-user System 10.5.204.43
Vsys 1, Flag 3
Port range: 20000 - 39999
Port size: start 200; max 2000
Block count 100, port count 20000
20000-20199: acme\administrator

Total host: 1
```

Send User Mappings to User-ID Using the XML API

User-ID provides many out-of-the box methods for obtaining user mapping information. However, you might have applications or devices that capture user information but cannot natively integrate with User-ID. For example, you might have a custom, internally developed application or a device that no standard user mapping method supports. In such cases, you can use the PAN-OS XML API to create custom scripts that send the information to the PAN-OS integrated User-ID agent or directly to the firewall. The PAN-OS XML API uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports POST and GET requests.

To enable an external system to send user mapping information to the PAN-OS integrated User-ID agent, create scripts that extract user login and logout events and use the events as input to the PAN-OS XML API request. Then define the mechanisms for submitting the XML API requests to the firewall (using cURL, for example) and use the API key of the firewall for secure communication. For more details, refer to the PAN-OS XML API Usage Guide.

© 2018 Palo Alto Networks, Inc.

Enable User- and Group-Based Policy

After you Enable User-ID, you will be able to configure Security Policy that applies to specific users and groups. User-based policy controls can also include application information (including which category and subcategory it belongs in, its underlying technology, or what the application characteristics are). You can define policy rules to safely enable applications based on users or groups of users, in either outbound or inbound directions.

Examples of user-based policies include:

- Enable only the IT department to use tools such as SSH, telnet, and FTP on standard ports.
- Allow the Help Desk Services group to use Slack.
- Allow all users to read Facebook, but block the use of Facebook apps, and restrict posting to employees in marketing.

© 2018 Palo Alto Networks, Inc.

Enable Policy for Users with Multiple Accounts

If a user in your organization has multiple responsibilities, that user might have multiple usernames (accounts), each with distinct privileges for accessing a particular set of services, but with all the usernames sharing the same IP address (the client system of the user). However, the User-ID agent can map any one IP address (or IP address and port range for terminal server users) to only one username for enforcing policy, and you can't predict which username the agent will map. To control access for all the usernames of a user, you must make adjustments to the rules, user groups, and User-ID agent.

For example, say the firewall has a rule that allows username corp_user to access email and a rule that allows username admin_user to access a MySQL server. The user logs in with either username from the same client IP address. If the User-ID agent maps the IP address to corp_user, then whether the user logs in as corp_user or admin_user, the firewall identifies that user as corp_user and allows access to email but not the MySQL server. On the other hand, if the User-ID agent maps the IP address to admin_user, the firewall always identifies the user as admin_user regardless of login and allows access to the MySQL server but not email. The following steps describe how to enforce both rules in this example.

STEP 1 | Configure a user group for each service that requires distinct access privileges.

In this example, each group is for a single service (email or MySQL server). However, it is common to configure each group for a set of services that require the same privileges (for example, one group for all basic user services and one group for all administrative services).

If your organization already has user groups that can access the services that the user requires, simply add the username that is used for less restricted services to those groups. In this example, the email server requires less restricted access than the MySQL server, and corp_user is the username for accessing email. Therefore, you add corp_user to a group that can access email (corp_employees) and to a group that can access the MySQL server (network_services).

If adding a username to a particular existing group would violate your organizational practices, you can create a custom group based on an LDAP filter. For this example, say network_services is a custom group, which you configure as follows:

1. Select **Device** > **User Identification** > **Group Mapping Settings** and **Add** a group mapping configuration with a unique **Name**.
2. Select an LDAP **Server Profile** and ensure the **Enabled** check box is enabled.
3. Select the **Custom Group** tab and **Add** a custom group with network_services as a **Name**.
4. Specify an **LDAP Filter** that matches an LDAP attribute of corp_user and click **OK**.
5. Click **OK** and **Commit**.



Later, if other users that are in the group for less restricted services are given additional usernames that access more restricted services, you can add those usernames to the group for more restricted services. This scenario is more common than the inverse; a user with access to more restricted services usually already has access to less restricted services.

STEP 2 | Configure the rules that control user access based on the groups you just configured.

Step 7 on page 454

1. Configure a security rule that allows the corp_employees group to access email.
2. Configure a security rule that allows the network_services group to access the MySQL server.

STEP 3 | Configure the ignore list of the User-ID agent.

© 2018 Palo Alto Networks, Inc.

This ensures that the User-ID agent maps the client IP address only to the username that is a member of the groups assigned to the rules you just configured. The ignore list must contain all the usernames of the user that are not members of those groups.

In this example, you add admin_user to the ignore list of the Windows-based User-ID agent to ensure that it maps the client IP address to corp_user. This guarantees that, whether the user logs in as corp_user or admin_user, the firewall identifies the user as corp_user and applies both rules that you configured because corp_user is a member of the groups that the rules reference.

1. Create an `ignore_user_list.txt` file.
2. Open the file and add admin_user.

   If you later add more usernames, each must be on a separate line.
3. Save the file to the User-ID agent folder on the domain server where the agent is installed.



If you use the PAN-OS integrated User-ID agent, see Configure User Mapping Using the PAN-OS Integrated User-ID Agent *on page 472* *for instructions on how to configure the ignore list.*

STEP 4 | Configure endpoint authentication for the restricted services.

This enables the endpoint to verify the credentials of the user and preserves the ability to enable access for users with multiple usernames.

In this example, you have configured a firewall rule that allows corp_user, as a member of the network_services group, to send a service request to the MySQL server. You must now configure the MySQL server to respond to any unauthorized username (such as corp_user) by prompting the user to enter the login credentials of an authorized username (admin_user).



If the user logs in to the network as admin_user, the user can then access the MySQL server without it prompting for the admin_user credentials again.

In this example, both corp_user and admin_user have email accounts, so the email server won't prompt for additional credentials regardless of which username the user entered when logging in to the network.

The firewall is now ready to enforce rules for a user with multiple usernames.

© 2018 Palo Alto Networks, Inc.

Verify the User-ID Configuration

After you configure user and group mapping, enable User-ID in your Security policy, and configure Authentication policy, you should verify that User-ID works properly.

STEP 1 | Access the firewall CLI.

STEP 2 | Verify that group mapping is working.

From the CLI, enter the following operational command:

```
> show user group-mapping statistics
```

STEP 3 | Verify that user mapping is working.

If you are using the PAN-OS integrated User-ID agent, you can verify this from the CLI using the following command:

```
> show user ip-user-mapping-mp all
IP              Vsys  From  User         Timeout (sec)
------------------------------------------------------
192.168.201.1   vsys1 UIA   acme\george            210
192.168.201.11  vsys1 UIA   acme\duane             210
192.168.201.50  vsys1 UIA   acme\betsy             210
192.168.201.10  vsys1 UIA   acme\administrator     210
192.168.201.100 vsys1 AD    acme\administrator     748
Total: 5 users
*: WMI probe succeeded
```

STEP 4 | Test your Security policy rule.

- From a machine in the zone where User-ID is enabled, attempt to access sites and applications to test the rules you defined in your policy and ensure that traffic is allowed and denied as expected.
- You can also use the **test security-policy-match** operational command to determine whether the policy is configured correctly. For example, suppose you have a rule that blocks user duane from playing World of Warcraft; you could test the policy as follows:

```
> test security-policy-match application worldofwarcraft source-user acme
\duane source any destination any destination-port any protocol 6
"deny worldofwarcraft" {
        from corporate;
        source any;
        source-region any;
        to internet;
        destination any;
        destination-region any;
        user acme\duane;
        category any;
        application/service worldofwarcraft;
        action deny;
        terminal no;
}
```

STEP 5 | Test your Authentication policy and Captive Portal configuration.

© 2018 Palo Alto Networks, Inc.

1. From the same zone, go to a machine that is not a member of your directory, such as a Mac OS system, and try to ping to a system external to the zone. The ping should work without requiring authentication.
2. From the same machine, open a browser and navigate to a web site in a destination zone that matches an Authentication rule you defined. The Captive Portal web form should display and prompt you for login credentials.
3. Log in using the correct credentials and confirm that you are redirected to the requested page.
4. You can also test your Authentication policy using the **test cp-policy-match** operational command as follows:

```
> test cp-policy-match from corporate to internet source 192.168.201.10
 destination 8.8.8.8
Matched rule: 'captive portal' action: web-form
```

STEP 6 | Verify that the log files display usernames.

Select a logs page (such as **Monitor** > **Logs** > **Traffic**) and verify that the Source User column displays usernames.

STEP 7 | Verify that reports display usernames.

1. Select **Monitor** > **Reports**.
2. Select a report type that includes usernames. For example, the Denied Applications report, Source User column, should display a list of the users who attempted to access the applications.

© 2018 Palo Alto Networks, Inc.

Deploy User-ID in a Large-Scale Network

A large-scale network can have hundreds of information sources that firewalls query to map IP addresses to usernames and to map usernames to user groups. You can simplify User-ID administration for such a network by aggregating the user mapping and group mapping information before the User-ID agents collect it, thereby reducing the number of required agents.

A large-scale network can also have numerous firewalls that use the mapping information to enforce policies. You can reduce the resources that the firewalls and information sources use in the querying process by configuring some firewalls to acquire mapping information through redistribution instead of direct querying. 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).

If you Configure Authentication Policy on page 212, your firewalls must also redistribute the Authentication Timestamps on page 211 associated with user responses to authentication challenges. Firewalls use the timestamps to evaluate the timeouts for Authentication policy rules. The timeouts allow a user who successfully authenticates to later request services and applications without authenticating again within the timeout periods. Redistributing timestamps enables you to enforce consistent timeouts for each user even if the firewall that initially grants a user access is not the same firewall that later controls access for that user.

- Deploy User-ID for Numerous Mapping Information Sources on page 501
- Redistribute User Mappings and Authentication Timestamps on page 505

Deploy User-ID for Numerous Mapping Information Sources

You can use Windows Log Forwarding and Global Catalog servers to simplify user mapping and group mapping in a large-scale network of Microsoft Active Directory (AD) domain controllers or Exchange servers. These methods simplify User-ID administration by aggregating the mapping information before the User-ID agents collect it, thereby reducing the number of required agents.

- Windows Log Forwarding and Global Catalog Servers on page 501
- Plan a Large-Scale User-ID Deployment on page 502
- Configure Windows Log Forwarding on page 503
- Configure User-ID for Numerous Mapping Information Sources on page 503

Windows Log Forwarding and Global Catalog Servers

Because each User-ID agent can monitor up to 100 servers, the firewall needs multiple User-ID agents to monitor a network with hundreds of AD domain controllers or Exchange servers. Creating and managing numerous User-ID agents involves considerable administrative overhead, especially in expanding networks where tracking new domain controllers is difficult. Windows Log Forwarding enables you to minimize the administrative overhead by reducing the number of servers to monitor and thereby reducing the number of User-ID agents to manage. When you configure Windows Log Forwarding, multiple domain controllers export their login events to a single domain member from which a User-ID agent collects the user mapping information.

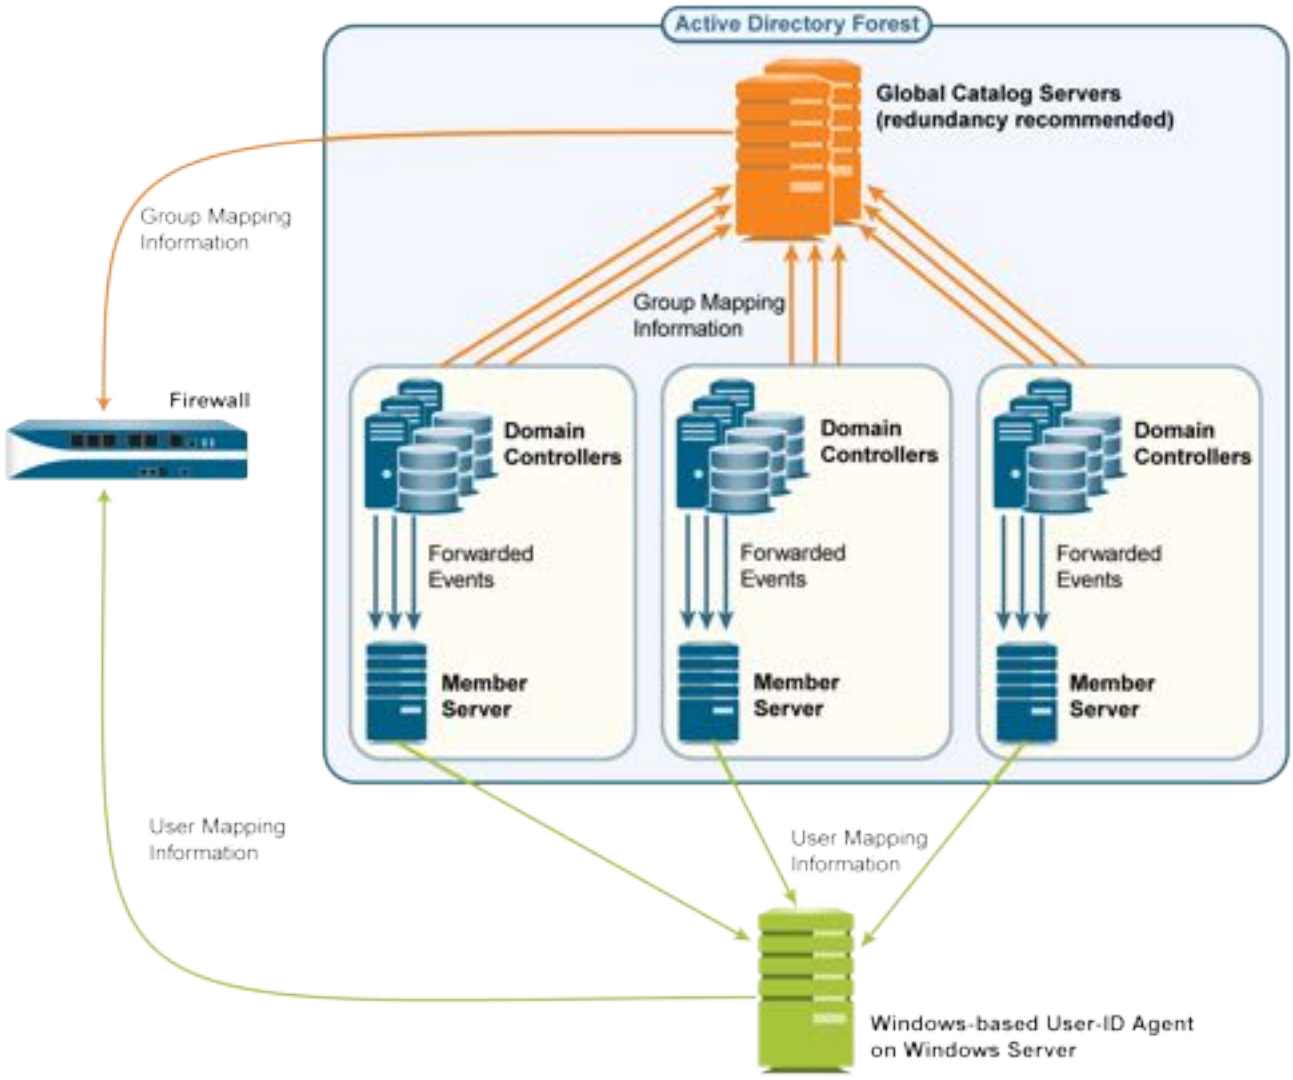


You can configure Windows Log Forwarding for Windows Server versions 2003, 2008, 2008 R2, 2012, and 2012 R2. Windows Log Forwarding is not available for non-Microsoft servers.

To collect group mapping information in a large-scale network, you can configure the firewall to query a Global Catalog server that receives account information from the domain controllers.

© 2018 Palo Alto Networks, Inc.

The following figure illustrates user mapping and group mapping for a large-scale network in which the firewall uses a Windows-based User-ID agent. See Plan a Large-Scale User-ID Deployment on page 502 to determine if this deployment suits your network.



Plan a Large-Scale User-ID Deployment

When deciding whether to use Windows Log Forwarding and Global Catalog servers for your User-ID implementation, consult your system administrator to determine:

- Bandwidth required for domain controllers to forward login events to member servers. The bandwidth is a multiple of the login rate (number of logins per minute) of the domain controllers and the byte size of each login event.

  Domain controllers won't forward their entire security logs; they forward only the events that the user mapping process requires per login: three events for Windows Server 2003 or four events for Windows Server 2008/2012 and MS Exchange.
- Whether the following network elements support the required bandwidth:
  - Domain controllers—Must support the processing load associated with forwarding the events.
  - Member Servers—Must support the processing load associated with receiving the events.

© 2018 Palo Alto Networks, Inc.

- Connections—The geographic distribution (local or remote) of the domain controllers, member servers, and Global Catalog servers is a factor. Generally, a remote distribution supports less bandwidth.

Configure Windows Log Forwarding

To configure Windows Log Forwarding, you need administrative privileges for configuring group policies on Windows servers. Configure Windows Log Forwarding on all the *Windows Event Collectors*—the member servers that collect login events from domain controllers. The following is an overview of the tasks; consult your Windows Server documentation for the specific steps.

STEP 1 | On each Windows Event Collector, enable event collection, add the domain controllers as event sources, and configure the event collection query (subscription). The events you specify in the subscription vary by domain controller platform:

- **Windows Server 2003**—The event IDs for the required events are 672 (Authentication Ticket Granted), 673 (Service Ticket Granted), and 674 (Ticket Granted Renewed).
- **Windows Server 2008/2012 (including R2) and 2016, or MS Exchange**—The event IDs for the required events are 4768 (Authentication Ticket Granted), 4769 (Service Ticket Granted), 4770 (Ticket Granted Renewed), and 4624 (Logon Success).



To forward events as quickly as possible, Minimize Latency when configuring the subscription.

User-ID agents monitor the Security log, not the default forwarded events location, on Windows Event Collectors. Therefore, perform the following steps on each Windows Event Collector to change the event logging path to the Security log.

1. Open the Event Viewer.
2. Right-click the **Security** log and select **Properties**.
3. Copy the **Log path** (default **`%SystemRoot%\System32\Winevt\Logs\security.evtx`**) and click **OK**.
4. Right-click the **Forwarded Events** folder and select **Properties**.
5. Replace the default **Log path** (**`%SystemRoot%\System32\Winevt\Logs\ForwardedEvents.evtx`**) by pasting the value from the **Security** log, and then click **OK**.

STEP 2 | Configure a group policy to enable Windows Remote Management (WinRM) on the domain controllers.

STEP 3 | Configure a group policy to enable Windows Event Forwarding on the domain controllers.

Configure User-ID for Numerous Mapping Information Sources

STEP 1 | Configure Windows Log Forwarding on the member servers that will collect login events.

Configure Windows Log Forwarding. This step requires administrative privileges for configuring group policies on Windows servers.

STEP 2 | Install the Windows-based User-ID agent.

Install the Windows-Based User-ID Agent on a Windows server that can access the member servers. Make sure the system that will host the User-ID agent is a member of the same domain as the servers it will monitor.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Configure the User-ID agent to collect user mapping information from the member servers.

1. Start the Windows-based User-ID agent.
2. Select **User Identification** > **Discovery** and perform the following steps for each member server that will receive events from domain controllers:
    1. In the Servers section, click **Add** and enter a **Name** to identify the member server.
    2. In the **Server Address** field, enter the FQDN or IP address of the member server.
    3. For the **Server Type**, select **Microsoft Active Directory**.
    4. Click **OK** to save the server entry.
3. Configure the remaining User-ID agent settings: see Configure the Windows-Based User-ID Agent for User Mapping.

STEP 4 | Configure an LDAP server profile to specify how the firewall connects to the Global Catalog servers (up to four) for group mapping information.



To improve availability, use at least two Global Catalog servers for redundancy.

You can collect group mapping information only for universal groups, not local domain groups (subdomains).

1. Select **Device** > **Server Profiles** > **LDAP**, click **Add**, and enter a **Name** for the profile.
2. In the Servers section, for each Global Catalog, click **Add** and enter the server **Name**, IP address (**LDAP Server**), and **Port**. For a plaintext or Start Transport Layer Security (Start TLS) connection, use **Port** 3268. For an LDAP over SSL connection, use **Port** 3269. If the connection will use Start TLS or LDAP over SSL, select the **Require SSL/TLS secured connection** check box.
3. In the **Base DN** field, enter the Distinguished Name (DN) of the point in the Global Catalog server where the firewall will start searching for group mapping information (for example, `DC=acbdomain,DC=com`).
4. For the **Type**, select **active-directory**.
5. Configure the remaining fields as necessary: see Add an LDAP server profile.

STEP 5 | Configure an LDAP server profile to specify how the firewall connects to the servers (up to four) that contain domain mapping information.

User-ID uses this information to map DNS domain names to NetBIOS domain names. This mapping ensures consistent domain/username references in policy rules.



To improve availability, use at least two servers for redundancy.

The steps are the same as for the LDAP server profile you created for Global Catalogs in the Step 4, except for the following fields:

- **LDAP Server**—Enter the IP address of the domain controller that contains the domain mapping information.
- **Port**—For a plaintext or Start TLS connection, use **Port** 389. For an LDAP over SSL connection, use **Port** 636. If the connection will use Start TLS or LDAP over SSL, select the **Require SSL/TLS secured connection** check box.
- **Base DN**—Select the DN of the point in the domain controller where the firewall will start searching for domain mapping information. The value must start with the string: `cn=partitions,cn=configuration` (for example, `cn=partitions,cn=configuration,DC=acbdomain,DC=com`).

© 2018 Palo Alto Networks, Inc.

STEP 6 | Create a group mapping configuration for each LDAP server profile you created.

1. Select **Device** > **User Identification** > **Group Mapping Settings**.
2. Click **Add** and enter a **Name** to identify the group mapping configuration.
3. Select the LDAP **Server Profile** and ensure the **Enabled** check box is selected.
4. Configure the remaining fields as necessary: see Map Users to Groups.

   If the Global Catalog and domain mapping servers reference more groups than your security rules require, configure the Group Include List and/or Custom Group list to limit the groups for which User-ID performs mapping.
5. Click **OK** and **Commit**.

Redistribute User Mappings and Authentication Timestamps

Every firewall that enforces user-based policy requires user mapping information. In a large-scale network, instead of configuring all your 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).

You can redistribute user mapping information collected through any method except Terminal Services (TS) agents. You cannot redistribute Group Mapping *or* HIP match *information.*

If you use Panorama and Dedicated Log Collectors to manage firewalls and aggregate firewall logs, you can use Panorama to manage User-ID redistribution*. Leveraging Panorama and your distributed log collection infrastructure is a simpler solution than creating extra connections between firewalls to redistribute User-ID information.*

If you Configure 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 Authentication policy rules. The timeouts allow a user who successfully authenticates to later request services and applications without authenticating again within the timeout periods. Redistributing 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 don't have to configure redistribution for each information type separately.

- Firewall Deployment for User-ID Redistribution
- Configure User-ID Redistribution

Firewall Deployment for User-ID Redistribution

To aggregate User-ID information, organize the redistribution sequence in layers, where each layer has one or more firewalls. In the bottom layer, PAN-OS integrated User-ID agents running on firewalls and Windows-based User-ID agents running on Windows servers map IP addresses to usernames. Each higher layer has firewalls that receive the mapping information and authentication timestamps from up to 100 redistribution points in the layer beneath it. The top-layer firewalls aggregate the mappings and timestamps from all layers. This deployment provides the option to configure policies for all users in top-layer firewalls and region- or function-specific policies for a subset of users in the corresponding domains served by lower-layer firewalls.

Figure 6: User-ID and Timestamp Redistribution on page 506 shows a deployment with three layers of firewalls that redistribute mappings and timestamps from local offices to regional offices and then to a global data center. The data center firewall that aggregates all the information shares it with other

© 2018 Palo Alto Networks, Inc.

data center firewalls so that they can all enforce policy and generate reports for users across your entire network. Only the bottom layer firewalls use User-ID agents to query the directory servers.

The information sources that the User-ID agents query do not count towards the maximum of ten *hops* in the sequence. However, Windows-based User-ID agents that forward mapping information to firewalls do count. Therefore, in this example, redistribution from the European region to all the data center firewalls requires only three hops, while redistribution from the North American region requires four hops. Also in this example, the top layer has two hops: the first to aggregate information in one data center firewall and the second to share the information with other data center firewalls.

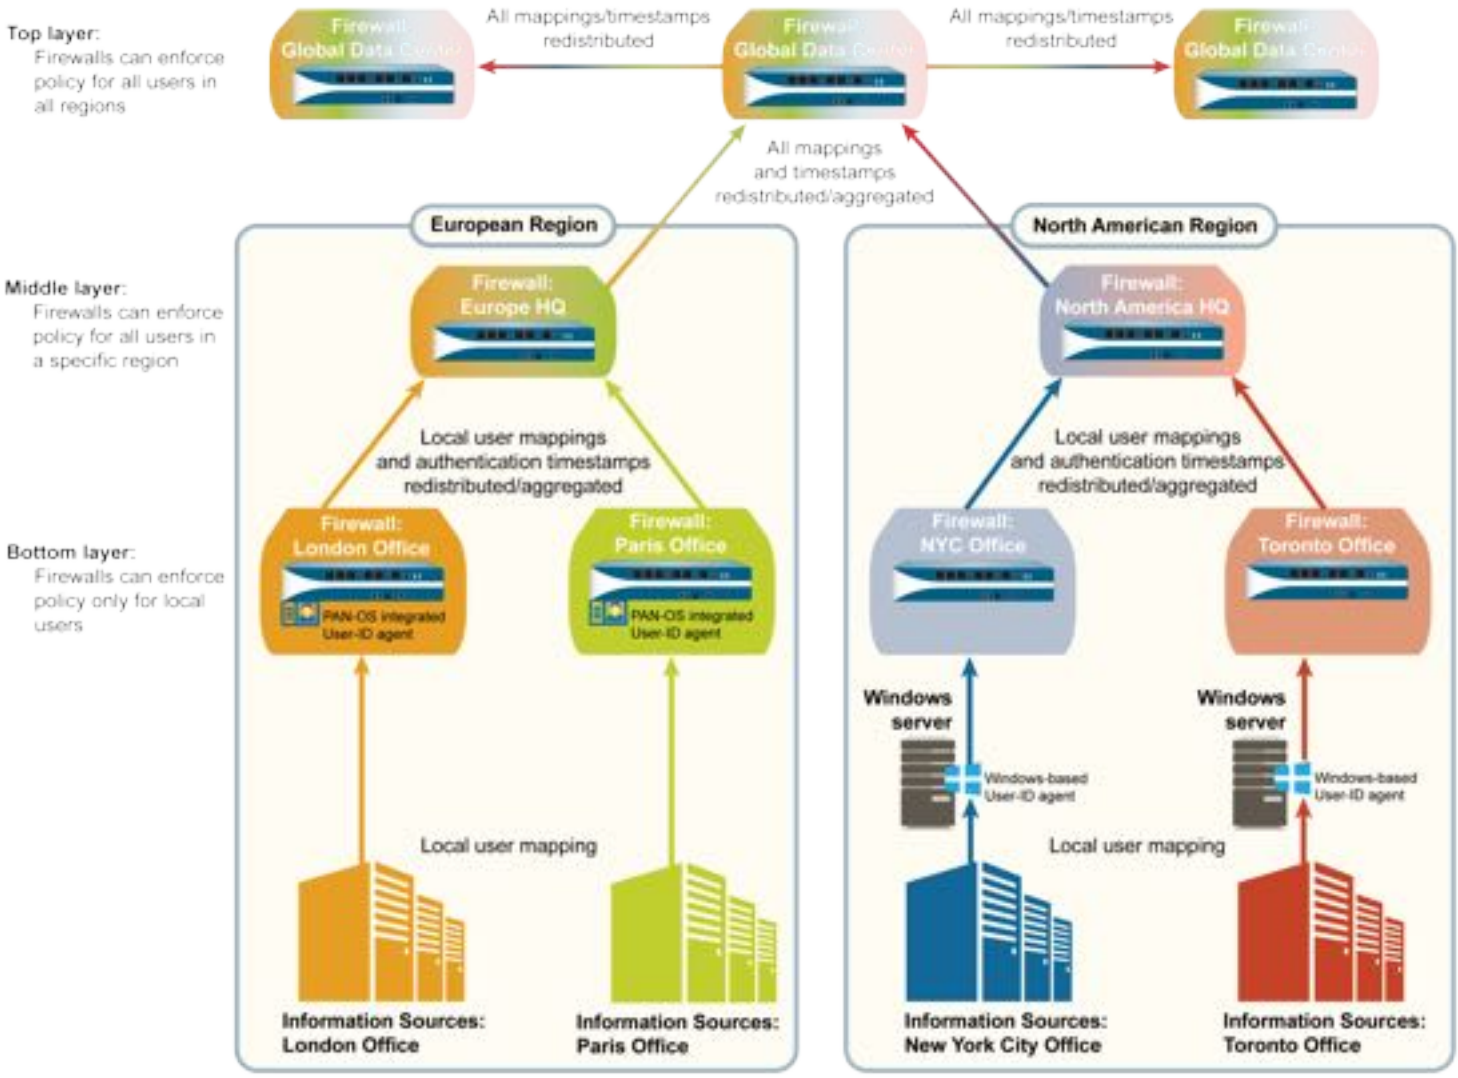


Figure 6: User-ID and Timestamp Redistribution

Configure User-ID Redistribution

Before you configure User-ID redistribution:

- ❑ Plan the redistribution architecture. Some factors to consider are:
  - Which firewalls will enforce policies for all users and which firewalls will enforce region- or function-specific policies for a subset of users?
  - How many hops does the redistribution sequence require to aggregate all User-ID information? The maximum allowed number of hops is ten.

© 2018 Palo Alto Networks, Inc.

- How can you minimize the number of firewalls that query the user mapping information sources? The fewer the number of querying firewalls, the lower the processing load is on both the firewalls and sources.

❑ Configure user mapping using PAN-OS Integrated User-ID agents or Windows-based User-ID agents.

❑ Configure Authentication Policy.

Perform the following steps on the firewalls in the User-ID redistribution sequence.

STEP 1 | Configure the firewall to redistribute User-ID information.

Skip this step if the firewall receives but does not redistribute User-ID information.

1. Select **Device** > **User Identification** > **User Mapping**.
2. (**Firewalls with multiple virtual systems only**) Select the **Location**. You must configure the User-ID settings for each virtual system.



You can redistribute information among virtual systems on different firewalls or on the same firewall. In both cases, each virtual system counts as one hop in the redistribution sequence.

3. Edit the Palo Alto Networks User-ID Agent Setup and select **Redistribution**.
4. Enter a **Collector Name** and **Pre-Shared Key** to identify this firewall or virtual system as a User-ID agent.
5. Click **OK** to save your changes.

STEP 2 | Configure the service route that the firewall uses to query other firewalls for User-ID information.

Skip this step if the firewall receives user mapping information from Windows-based User-ID agents or directly from the information sources (such as directory servers) instead of from other firewalls.

1. Select **Device** > **Setup** > **Services**.
2. (**Firewalls with multiple virtual systems only**) Select **Global** (for a firewall-wide service route) or **Virtual Systems** (for a virtual system-specific service route), and then configure the service route.
3. Click **Service Route Configuration**, select **Customize**, and select **IPv4** or **IPv6** based on your network protocols. Configure the service route for both protocols if your network uses both.
4. Select **UID Agent** and then select the **Source Interface** and **Source Address**.
5. Click **OK** twice to save the service route.

STEP 3 | Enable the firewall to respond when other firewalls query it for User-ID information.

Skip this step if the firewall receives but does not redistribute User-ID information.

Configure an Interface Management Profile with the **User-ID** service enabled and assign the profile to a firewall interface.

STEP 4 | Commit and verify your changes.

1. **Commit** your changes to activate them.
2. Access the CLI of a firewall that redistributes User-ID information.
3. Display all the user mappings by running the following command:

```
> show user ip-user-mapping all
```

4. Record the IP address associated with any username.
5. Access the CLI of a firewall that receives redistributed User-ID information.
6. Display the mapping information and authentication timestamp for the *<ip_address>* you recorded:

© 2018 Palo Alto Networks, Inc.

```
> show user ip-user-mapping <ip_address>
IP address:     192.0.2.0 (vsys1)
User:           corpdomain\username1
From:           UIA
Idle Timeout:   10229s
Max. TTL:       10229s
MFA Timestamp: first(1) - 2016/12/09 08:35:04
Group(s):       corpdomain\groupname(621)
```



This example output shows the authentication timestamp for one response to an authentication challenge (factor). For Authentication policy rules that use Multi-Factor Authentication *(MFA), the output shows multiple* Authentication Timestamps*.*

App-ID

To safely enable applications on your network, the Palo Alto Networks next-generation firewalls provide both an application and web perspective—App-ID and URL Filtering—to protect against a full spectrum of legal, regulatory, productivity, and resource utilization risks.

App-ID enables visibility into the applications on the network, so you can learn how they work and understand their behavioral characteristics and their relative risk. This application knowledge allows you to create and enforce security policy rules to enable, inspect, and shape desired applications and block unwanted applications. When you define policy rules to allow traffic, App-ID begins to classify traffic without any additional configuration.

- App-ID Overview on page 511
- Manage Custom or Unknown Applications on page 512
- Manage New App-IDs Introduced in Content Releases on page 513
- Use Application Objects in Policy on page 518
- Applications with Implicit Support on page 523
- Application Level Gateways on page 526
- Disable the SIP Application-level Gateway (ALG) on page 528

© 2018 Palo Alto Networks, Inc.

App-ID Overview

App-ID, a patented traffic classification system only available in Palo Alto Networks firewalls, determines what an application is irrespective of port, protocol, encryption (SSH or SSL) or any other evasive tactic used by the application. It applies multiple classification mechanisms—application signatures, application protocol decoding, and heuristics—to your network traffic stream to accurately identify applications.

Here's how App-ID identifies applications traversing your network:

- Traffic is matched against policy to check whether it is allowed on the network.
- Signatures are then applied to allowed traffic to identify the application based on unique application properties and related transaction characteristics. The signature also determines if the application is being used on its default port or it is using a non-standard port. If the traffic is allowed by policy, the traffic is then scanned for threats and further analyzed for identifying the application more granularly.
- If App-ID determines that encryption (SSL or SSH) is in use, and a Decryption on page 589 policy rule is in place, the session is decrypted and application signatures are applied again on the decrypted flow.
- Decoders for known protocols are then used to apply additional context-based signatures to detect other applications that may be tunneling inside of the protocol (for example, Yahoo! Instant Messenger used across HTTP). Decoders validate that the traffic conforms to the protocol specification and provide support for NAT traversal and opening dynamic pinholes for applications such as SIP and FTP.
- For applications that are particularly evasive and cannot be identified through advanced signature and protocol analysis, heuristics or behavioral analysis may be used to determine the identity of the application.

When the application is identified, the policy check determines how to treat the application, for example—block, or allow and scan for threats, inspect for unauthorized file transfer and data patterns, or shape using QoS.

© 2018 Palo Alto Networks, Inc.

Manage Custom or Unknown Applications

Palo Alto Networks provides weekly application updates to identify new App-ID signatures. By default, App-ID is always enabled on the firewall, and you don't need to enable a series of signatures to identify well-known applications. Typically, the only applications that are classified as unknown traffic—tcp, udp or non-syn-tcp—in the ACC and the traffic logs are commercially available applications that have not yet been added to App-ID, internal or custom applications on your network, or potential threats.

On occasion, the firewall may report an application as unknown for the following reasons:

- Incomplete data—A handshake took place, but no data packets were sent prior to the timeout.
- Insufficient data—A handshake took place followed by one or more data packets; however, not enough data packets were exchanged to identify the application.

The following choices are available to handle unknown applications:

- Create security policies to control unknown applications by unknown TCP, unknown UDP or by a combination of source zone, destination zone, and IP addresses.
- Request an App-ID from Palo Alto Networks—If you would like to inspect and control the applications that traverse your network, for any unknown traffic, you can record a packet capture. If the packet capture reveals that the application is a commercial application, you can submit this packet capture to Palo Alto Networks for App-ID development. If it is an internal application, you can create a custom App-ID and/or define an application override policy.
- Create a Custom Application on page 519 with a signature and attach it to a security policy, or create a custom application and define an application override policy—A custom application allows you to customize the definition of the internal application—its characteristics, category and sub-category, risk, port, timeout—and exercise granular policy control in order to minimize the range of unidentified traffic on your network. Creating a custom application also allows you to correctly identify the application in the **ACC** and traffic logs and is useful in auditing/reporting on the applications on your network. For a custom application you can specify a signature and a pattern that uniquely identifies the application and attach it to a security policy that allows or denies the application.

  Alternatively, if you would like the firewall to process the custom application using fast path (Layer-4 inspection instead of using App-ID for Layer-7 inspection), you can reference the custom application in an application override policy rule. An application override with a custom application will prevent the session from being processed by the App-ID engine, which is a Layer-7 inspection. Instead it forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4, and thereby saves application processing time.

  For example, if you build a custom application that triggers on a host header *www.mywebsite.com*, the packets are first identified as *web-browsing* and then are matched as your custom application (whose parent application is web-browsing). Because the parent application is web-browsing, the custom application is inspected at Layer-7 and scanned for content and vulnerabilities.

  If you define an application override, the firewall stops processing at Layer-4. The custom application name is assigned to the session to help identify it in the logs, and the traffic is not scanned for threats.

© 2018 Palo Alto Networks, Inc.

Manage New App-IDs Introduced in Content Releases

Installing new App-IDs included in a content release version can sometimes cause a change in policy enforcement for the now uniquely-identified application. Before installing a new content release, review the policy impact for new App-IDs and stage any necessary policy updates. Assess the treatment an application receives both before and after the new content is installed. You can then modify existing security policy rules using the new App-IDs contained in a downloaded content release (prior to installing the App-IDs). This enables you to simultaneously update your security policies and install new content, and allows for a seamless shift in policy enforcement. Alternatively, you can also choose to disable new App-IDs when installing a new content release version; this enables protection against the latest threats, while giving you the flexibility to enable the new App-IDs after you've had the chance to prepare any policy changes.

The following options enable you to assess the impact of new App-IDs on existing policy enforcement, disable (and enable) App-IDs, and seamlessly update policy rules to secure and enforce newly-identified applications:

- Review New App-IDs
- Disable or Enable App-IDs
- Prepare Policy Updates for Pending App-IDs

Review New App-IDs

Review new App-ID signatures introduced in a Applications and/or Threats content update. For each new application signature introduced, you can preview the App-ID details, including a description of the application identified by the App-ID, other existing App-IDs that the new signature is dependent on (such as SSL or HTTP), and the category the application traffic received before the introduction of the new App-ID (for example, an application might be classified as web-browsing traffic before a App-ID signature is introduced that uniquely identifies the traffic). After reviewing the description and details for a new App-ID signature, review the App-ID signature impact on existing policy enforcement. When new application signatures are introduced, the newly-identified application traffic might no longer match to policies that previously enforced the application. Reviewing the policy impact for new application signatures enables you to identify the policies that will no longer enforce the application when the new App-ID is installed.

After downloading a new content release version, review the new App-IDs included in the content version and assess the impact of the new App-IDs on existing policy rules:

- Review New App-IDs Since Last Content Version on page 513
- Review New App-ID Impact on Existing Policy Rules on page 514

Review New App-IDs Since Last Content Version

STEP 1 | Select **Device** > **Dynamic Updates** and select **Check Now** to refresh the list of available content updates.

STEP 2 | **Download** the latest Applications and Threats content update. When the content update is downloaded, an **Apps** link will appear in the Features column for that content update.

STEP 3 | Click the **Apps** link in the **Features** column to view details on newly-identified applications:

© 2018 Palo Alto Networks, Inc.



A list of App-IDs shows all new App-IDs introduced from the content version installed on the firewall, to the selected **Content Version**.

App-ID details that you can use to assess possible impact to policy enforcement include:

- **Depends on**—Lists the application signatures that this App-ID relies on to uniquely identify the application. If one of the application signatures listed in the **Depends On** field is disabled, the dependent App-ID is also disabled.
- **Previously Identified As**—Lists the App-IDs that matched to the application before the new App-ID was installed to uniquely identify the application.
- **App-ID Enabled**—All App-IDs display as enabled when a content release is downloaded, unless you choose to manually disable the App-ID signature before installing the content update (see Disable or Enable App-IDs on page 515).

  Multi-vsys firewalls display App-ID status as **vsys-specific**. This is because the status is not applied across virtual systems and must be individually enabled or disabled for each virtual system. To view the App-ID status for a specific virtual system, select **Objects** > **Applications**, select a **Virtual System**, and select the App-ID.



STEP 4 | Next steps...

- Disable or Enable App-IDs on page 515.
- Prepare Policy Updates for Pending App-IDs on page 516.

Review New App-ID Impact on Existing Policy Rules

STEP 1 | Select **Device** > **Dynamic Updates**.

STEP 2 | You can review the policy impact of new content release versions that are downloaded to the firewall. **Download** a new content release version, and click the **Review Policies** in the Action column. The **Policy review based on candidate configuration** dialog allows you to filter by **Content Version** and view App-IDs introduced in a specific release (you can also filter the policy impact of new App-IDs according to **Rulebase** and **Virtual System**).

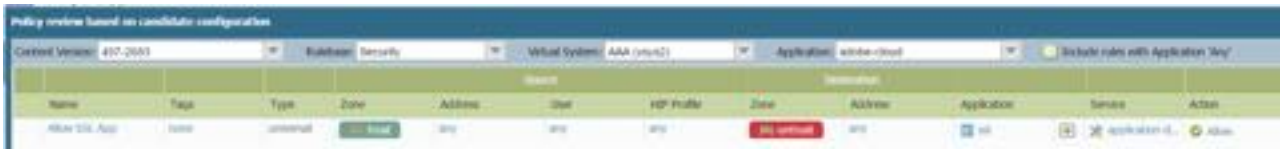
© 2018 Palo Alto Networks, Inc.

STEP 3 | Select a new App-ID from the **Application** drop-down to view policy rules that currently enforce the application. The rules displayed are based on the applications signatures that match to the application before the new App-ID is installed (view application details to see the list of application signatures that an application was **Previously Identified As** before the new App-ID).

STEP 4 | Use the detail provided in the policy review to plan policy rule updates to take effect when the App-ID is installed and enabled to uniquely identify the application.

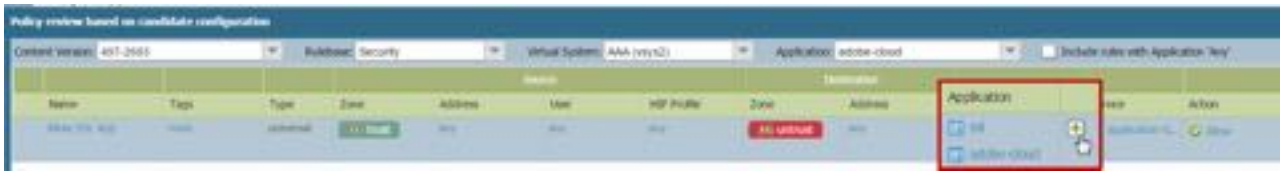
You can continue to Prepare Policy Updates for Pending App-IDs on page 516, or you can directly add the new App-ID to policy rules that the application was previously matched to by continuing to use the policy review dialog.

In the following example, the new App-ID adobe-cloud is introduced in a content release. Adobe-cloud traffic is currently identified as SSL and web-browsing traffic. Policy rules configured to enforce SSL or web-browsing traffic are listed to show what policy rules will be affected when the new App-ID is installed. In this example, the rule Allow SSL App currently enforces SSL traffic. To continue to allow adobe-cloud traffic when it is uniquely identified, and no longer identified as SSL traffic.



Add ⊞ the new App-ID to existing policy rules, to allow the application traffic to continue to be enforced according to your existing security requirements when the App-ID is installed.

In this example, to continue to allow adobe-cloud traffic when it is uniquely identified by the new App-ID, and no longer identified as SSL traffic, add the new App-ID to the security policy rule Allow SSL App.



The policy rule updates take effect only when the application updates are installed.

STEP 5 | Next steps:

- Disable or Enable App-IDs on page 515.
- Prepare Policy Updates for Pending App-IDs on page 516.

Disable or Enable App-IDs

Disable new App-IDs included in a content release to immediately benefit from protection against the latest threats while continuing to have the flexibility to later enable App-IDs after preparing necessary policy updates. You can disable all App-IDs introduced in a content release, set scheduled content updates to automatically disable new App-IDs, or disable App-IDs for specific applications.

Policy rules referencing App-IDs only match to and enforce traffic based on enabled App-IDs.

Certain App-IDs cannot be disabled and only allow a status of enabled. App-IDs that cannot be disabled included some application signatures implicitly used by other App-IDs (such as unknown-tcp). Disabling a base App-ID could cause App-IDs which depend on the base App-ID to also be disabled. For example, disabling facebook-base will disable all other Facebook App-IDs.

© 2018 Palo Alto Networks, Inc.

- Disable all App-IDs in a content release or for scheduled content updates.
  - To disable all new App-IDs introduced in a content release, select **Device** > **Dynamic Updates** and **Install** an Application and Threats content release. When prompted, select **Disable new apps in content update**. Select the check box to disable apps and continue installing the content update; this allows you to be protected against threats, and gives you the option to enable the apps at a later time.
  - On the **Device** > **Dynamic Updates** page, select **Schedule**. Choose to **Disable new apps in content update** for downloads and installations of content releases.

- Disable App-IDs for one application or multiple applications at a single time.
  - To quickly disable a single application or multiple applications at the same time, click **Objects** > **Applications**. Select one or more application check box and click **Disable**.
  - To review details for a single application, and then disable the App-ID for that application, select **Objects** > **Applications** and **Disable App-ID**. You can use this step to disable both pending App-IDs (where the content release including the App-ID is downloaded to the firewall but not installed) or installed App-IDs.

- Enable App-IDs.

  Enable App-IDs that you previously disabled by selecting **Objects** > **Applications**. Select one or more application check box and click **Enable** or open the details for a specific application and click **Enable App-ID**.

Prepare Policy Updates for Pending App-IDs

You can now stage seamless policy updates for new App-IDs. Release versions prior to PAN-OS 7.0 required you to install new App-IDs (as part of a content release) and then make necessary policy updates. This allowed for a period during which the newly-identified application traffic was not enforced, either by existing rules (that the traffic had matched to before being uniquely identified) or by rules that had yet to be created or modified to use the new App-ID.

Pending App-IDs can now be added to policy rules to prevent gaps in policy enforcement that could occur during the period between installing a content release and updating security policy. Pending App-IDs includes App-IDs that have been manually disabled, or App-IDs that are downloaded to the firewall but not installed. Pending App-IDs can be used to update policies both before and after installing a new content release. Though they can be added to policy rules, pending App-IDs are not enforced until the App-IDs are both installed and enabled on the firewall.

The names of App-IDs that have been manually disabled display as gray and italicized, to indicate the disabled status:

- Disabled App-ID listed on the **Objects** > **Applications** page:



- Disabled App-ID included in a security policy rule:



© 2018 Palo Alto Networks, Inc.

App-IDs that are included in a downloaded content release version might have an App-ID status of enabled, but App-IDs are not enforced until the corresponding content release version is installed.

- To install the content release version now and then update policies:

  Do this to benefit from new threat signatures immediately, while you review new application signatures and update your policies.

  1. Select **Device** > **Dynamic Updates** and **Download** the latest content release version.
  2. Review New App-ID Impact on Existing Policy Rules to assess the policy impact of new App-IDs.
  3. **Install** the latest content release version. Before the content release is installed, you are prompted to **Disable new apps in content update**. Select the check box and continue to install the content release. Threat signatures included in the content release will be installed and effective, while new or updated App-IDs are disabled.
  4. Select **Policies** and update **Security**, **QoS**, and **Policy Based Forwarding** rules to match to and enforce the now uniquely identified application traffic, using the pending App-IDs.
  5. Select **Objects** > **Applications** and select one or multiple disabled App-IDs and click **Enable**.
  6. **Commit** your changes to seamlessly update policy enforcement for new App-IDs.

- Update policies now and then install the content release version.
  1. Select **Device** > **Dynamic Updates** and **Download** the latest content release version.
  2. Review New App-ID Impact on Existing Policy Rules to assess the policy impact of new App-IDs.
  3. While reviewing the policy impact for new App-IDs, you can use the **Policy Review based on candidate configuration** to add a new App-ID to existing policy rules: ⊞. The new App-ID is added to the existing rules as a disabled App-ID.
  4. Continue to review the policy impact for all App-IDs included in the latest content release version by selecting App-IDs in the **Applications** drop-down. Add the new App-IDs to existing policies as needed. Click **OK** to save your changes.
  5. **Install** the latest content release version.
  6. **Commit** your changes to seamlessly update policy enforcement for new App-IDs.

© 2018 Palo Alto Networks, Inc.

Use Application Objects in Policy

- Create an Application Group on page 518
- Create an Application Filter on page 518
- Create a Custom Application on page 519

Create an Application Group

An application group is an object that contains applications that you want to treat similarly in policy. Application groups are useful for enabling access to applications that you explicitly sanction for use within your organization. Grouping sanctioned applications simplifies administration of your rulebases. Instead of having to update individual policy rules when there is a change in the applications you support, you can update only the affected application groups.

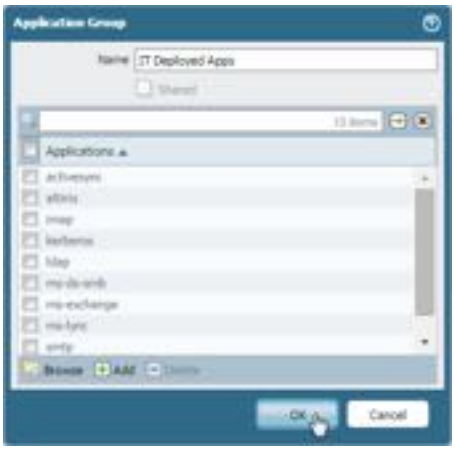
When deciding how to group applications, consider how you plan to enforce access to your sanctioned applications and create an application group that aligns with each of your policy goals. For example, you might have some applications that you will only allow your IT administrators to access, and other applications that you want to make available for any known user in your organization. In this case, you would create separate application groups for each of these policy goals. Although you generally want to enable access to applications on the default port only, you may want to group applications that are an exception to this and enforce access to those applications in a separate rule.

STEP 1 | Select **Objects** > **Application Groups**.

STEP 2 | **Add** a group and give it a descriptive **Name**.

STEP 3 | (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.

STEP 4 | **Add** the applications you want in the group and then click **OK**.



STEP 5 | **Commit** the configuration.

Create an Application Filter

An application filter is an object that dynamically groups applications based on application attributes that you define, including category, subcategory, technology, risk factor, and characteristic. This is useful when

© 2018 Palo Alto Networks, Inc.

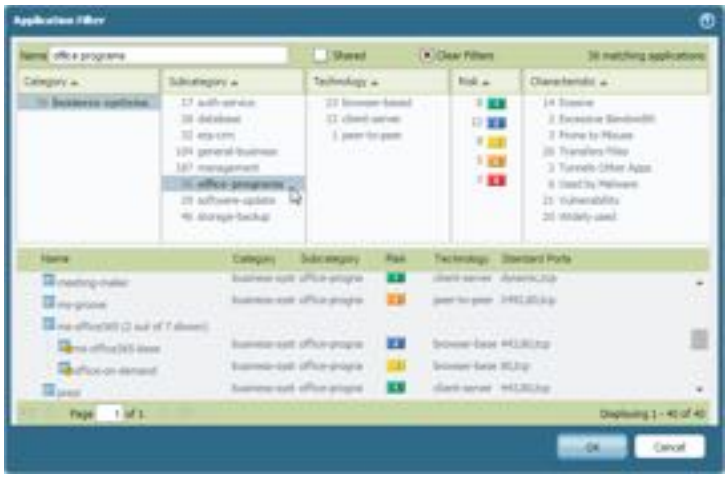
you want to safely enable access to applications that you do not explicitly sanction, but that you want users to be able to access. For example, you may want to enable employees to choose their own office programs (such as Evernote, Google Docs, or Microsoft Office 365) for business use. To safely enable these types of applications, you could create an application filter that matches on the Category **business-systems** and the Subcategory **office-programs**. As new applications office programs emerge and new App-IDs get created, these new applications will automatically match the filter you defined; you will not have to make any additional changes to your policy rulebase to safely enable any application that matches the attributes you defined for the filter.

STEP 1 | Select **Objects** > **Application Filters**.

STEP 2 | **Add** a filter and give it a descriptive **Name**.

STEP 3 | (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.

STEP 4 | Define the filter by selecting attribute values from the Category, Subcategory, Technology, Risk, and Characteristic sections. As you select values, notice that the list of matching applications at the bottom of the dialog narrows. When you have adjusted the filter attributes to match the types of applications you want to safely enable, click **OK**.



STEP 5 | **Commit** the configuration.

Create a Custom Application

To safely enable applications you must classify all traffic, across all ports, all the time. With App-ID, the only applications that are typically classified as unknown traffic—tcp, udp or non-syn-tcp—in the ACC and the Traffic logs are commercially available applications that have not yet been added to App-ID, internal or custom applications on your network, or potential threats.

If you are seeing unknown traffic for a commercial application that does not yet have an App-ID, you can submit a request for a new App-ID here: http://researchcenter.paloaltonetworks.com/submit-an-application/.

To ensure that your internal custom applications do not show up as unknown traffic, create a custom application. You can then exercise granular policy control over these applications in order to minimize the range of unidentified traffic on your network, thereby reducing the attack surface. Creating a custom

© 2018 Palo Alto Networks, Inc.

application also allows you to correctly identify the application in the ACC and Traffic logs, which enables you to audit/report on the applications on your network.

To create a custom application, you must define the application attributes: its characteristics, category and sub-category, risk, port, timeout. In addition, you must define patterns or values that the firewall can use to match to the traffic flows themselves (the *signature*). Finally, you can attach the custom application to a security policy that allows or denies the application (or add it to an application group or match it to an application filter). You can also create custom applications to identify ephemeral applications with topical interest, such as ESPN3-Video for world cup soccer or March Madness.



In order to collect the right data to create a custom application signature, you'll need a good understanding of packet captures and how datagrams are formed. If the signature is created too broadly, you might inadvertently include other similar traffic; if it is defined too narrowly, the traffic will evade detection if it does not strictly match the pattern.

Custom applications are stored in a separate database on the firewall and this database is not impacted by the weekly App-ID updates.

The supported application protocol decoders that enable the firewall to detect applications that may be tunneling inside of the protocol include the following as of content release version 609: FTP, HTTP, IMAP, POP3, SMB, and SMTP.

The following is a basic example of how to create a custom application.

STEP 1 | Gather information about the application that you will be able to use to write custom signatures.

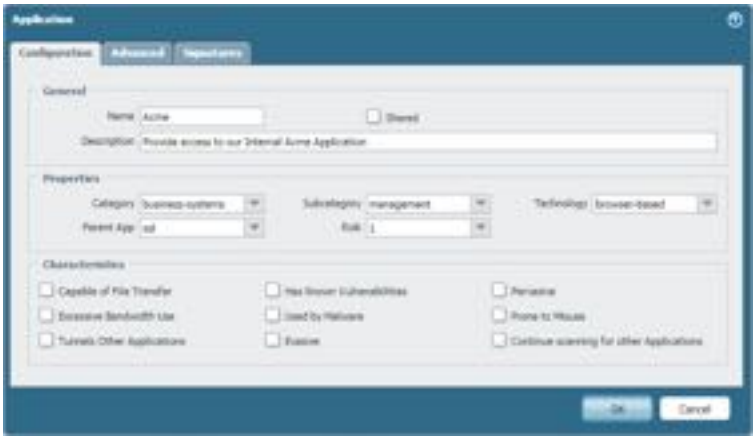
To do this, you must have an understanding of the application and how you want to control access to it. For example, you may want to limit what operations users can perform within the application (such as uploading, downloading, or live streaming). Or you may want to allow the application, but enforce QoS policing.

- Capture application packets so that you can find unique characteristics about the application on which to base your custom application signature. One way to do this is to run a protocol analyzer, such as Wireshark, on the client system to capture the packets between the client and the server. Perform different actions in the application, such as uploading and downloading, so that you will be able to locate each type of session in the resulting packet captures (PCAPs).
- Because the firewall by default takes packet captures for all unknown traffic, if the firewall is between the client and the server you can view the packet capture for the unknown traffic directly from the Traffic log.
- Use the packet captures to find patterns or values in the packet *contexts* that you can use to create signatures that will uniquely match the application traffic. For example, look for string patterns in HTTP response or request headers, URI paths, or hostnames. For information on the different string contexts you can use to create application signatures and where you can find the corresponding values in the packet, refer to Creating Custom Threat Signatures.

STEP 2 | Add the custom application.

1. Select **Objects** > **Applications** and click **Add**.
2. On the **Configuration** tab, enter a **Name** and a **Description** for the custom application that will help other administrators understand why you created the application.
3. (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.
4. Define the application Properties and Characteristics.

© 2018 Palo Alto Networks, Inc.

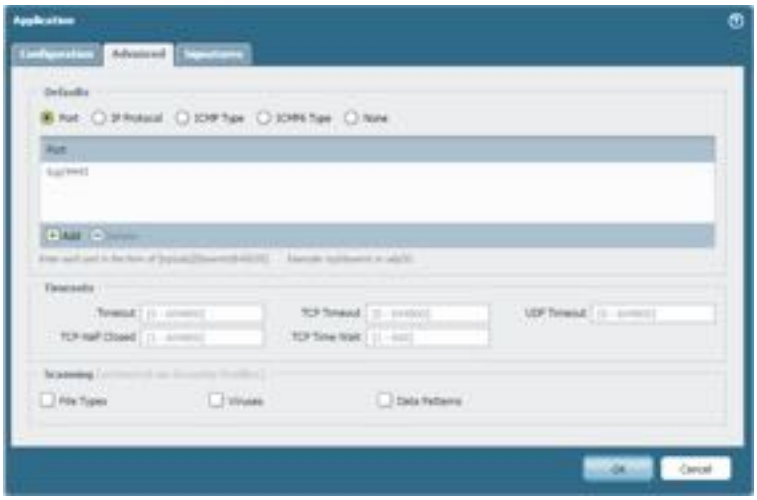


STEP 3 | Define details about the application, such as the underlying protocol, the port number the application runs on, the timeout values, and any types of scanning you want to be able to perform on the traffic.

On the **Advanced** tab, define settings that will allow the firewall to identify the application protocol:

- Specify the default ports or protocol that the application uses.
- Specify the session timeout values. If you don't specify timeout values, the default timeout values will be used.
- Indicate any type of additional scanning you plan to perform on the application traffic.

For example, to create a custom TCP-based application that runs over SSL, but uses port 4443 (instead of the default port for SSL, 443), you would specify the port number. By adding the port number for a custom application, you can create policy rules that use the default port for the application rather than opening up additional ports on the firewall. This improves your security posture.



STEP 4 | Define the criteria that the firewall will use to match the traffic to the new application.

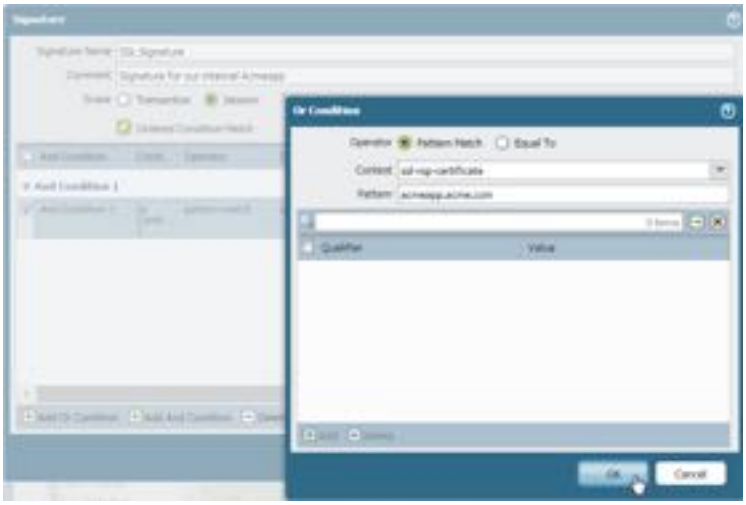
You will use the information you gathered from the packet captures to specify unique string context values that the firewall can use to match patterns in the application traffic.

1. On the **Signatures** tab, click **Add** and define a **Signature Name** and optionally a **Comment** to provide information about how you intend to use this signature.
2. Specify the **Scope** of the signature: whether it matches to a full **Session** or a single **Transaction**.
3. Specify conditions to define signatures by clicking **Add And Condition** or **Add Or Condition**.

© 2018 Palo Alto Networks, Inc.

4. Select an **Operator** to define the type of match conditions you will use: **Pattern Match** or **Equal To**.
   - If you selected **Pattern Match**, select the **Context** and then use a regular expression to define the **Pattern** to match the selected context. Optionally, click **Add** to define a qualifier/value pair. The **Qualifier** list is specific to the **Context** you chose.
   - If you selected **Equal To**, select the **Context** and then use a regular expression to define the **Position** of the bytes in the packet header to use match the selected context. Choose from **first-4bytes** or **second-4bytes**. Define the 4-byte hex value for the **Mask** (for example, 0xffffff00) and **Value** (for example, 0xaabbccdd).

   For example, if you are creating a custom application for one of your internal applications, you could use the **ssl-rsp-certificate Context** to define a pattern match for the certificate response message of a SSL negotiation from the server and create a **Pattern** to match the commonName of the server in the message as shown here:



5. Repeat steps c and d for each matching condition.
6. If the order in which the firewall attempts to match the signature definitions is important, make sure the **Ordered Condition Match** check box is selected and then order the conditions so that they are evaluated in the appropriate order. Select a condition or a group and click **Move Up** or **Move Down**. You cannot move conditions from one group to another.
7. Click **OK** to save the signature definition.

STEP 5 | Save the application.

1. Click **OK** to save the custom application definition.
2. Click **Commit**.

STEP 6 | Validate that traffic matches the custom application as expected.

1. Select **Policies** > **Security** and **Add** a security policy rule to allow the new application.
2. Run the application from a client system that is between the firewall and the application and then check the Traffic logs (**Monitor** > **Traffic**) to make sure that you see traffic matching the new application (and that it is being handled per your policy rule).

© 2018 Palo Alto Networks, Inc.

Applications with Implicit Support

When creating a policy to allow specific applications, you must also be sure that you are allowing any other applications on which the application depends. In many cases, you do not have to explicitly allow access to the dependent applications in order for the traffic to flow because the firewall is able to determine the dependencies and allow them implicitly. This implicit support also applies to custom applications that are based on HTTP, SSL, MS-RPC, or RTSP. Applications for which the firewall cannot determine dependent applications on time will require that you explicitly allow the dependent applications when defining your policies. You can determine application dependencies in Applipedia.

The following table lists the applications for which the firewall has implicit support (as of Content Update 595).

Table 2: Applications with Implicit Support

Application	Implicitly Supports
360-safeguard-update	http
apple-update	http
apt-get	http
as2	http
avg-update	http
avira-antivir-update	http, ssl
blokus	rtmp
bugzilla	http
clubcooee	http
corba	http
cubby	http, ssl
dropbox	ssl
esignal	http
evernote	http, ssl
ezhelp	http
facebook	http, ssl
facebook-chat	jabber
facebook-social-plugin	http

© 2018 Palo Alto Networks, Inc.

Application	Implicitly Supports
fastviewer	http, ssl
forticlient-update	http
good-for-enterprise	http, ssl
google-cloud-print	http, ssl, jabber
google-desktop	http
google-talk	jabber
google-update	http
gotomypc-desktop-sharing	citrix-jedi
gotomypc-file-transfer	citrix-jedi
gotomypc-printing	citrix-jedi
hipchat	http
iheartradio	ssl, http, rtmp
infront	http
instagram	http, ssl
issuu	http, ssl
java-update	http
jepptech-updates	http
kerberos	rpc
kik	http, ssl
lastpass	http, ssl
logmein	http, ssl
mcafee-update	http
megaupload	http
metatrader	http
mocha-rdp	t_120
mount	rpc

© 2018 Palo Alto Networks, Inc.

Application	Implicitly Supports
ms-frs	msrpc
ms-rdp	t_120
ms-scheduler	msrpc
ms-service-controller	msrpc
nfs	rpc
oovoo	http, ssl
paloalto-updates	ssl
panos-global-protect	http
panos-web-interface	http
pastebin	http
pastebin-posting	http
pinterest	http, ssl
portmapper	rpc
prezi	http, ssl
rdp2tcp	t_120
renren-im	jabber
roboform	http, ssl
salesforce	http
stumbleupon	http
supremo	http
symantec-av-update	http
trendmicro	http
trillian	http, ssl
twitter	http
whatsapp	http, ssl
xm-radio	rtsp

© 2018 Palo Alto Networks, Inc.

Application Level Gateways

The Palo Alto Networks firewall does not classify traffic by port and protocol; instead it identifies the application based on its unique properties and transaction characteristics using the App-ID technology. Some applications, however, require the firewall to dynamically open *pinholes* to establish the connection, determine the parameters for the session and negotiate the ports that will be used for the transfer of data; these applications use the application-layer payload to communicate the dynamic TCP or UDP ports on which the application opens data connections. For such applications, the firewall serves as an Application Level Gateway (ALG), and it opens a pinhole for a limited time and for exclusively transferring data or control traffic. The firewall also performs a NAT rewrite of the payload when necessary.



- *H.323 (H.225 and H.248) ALG is not supported in gatekeeper routed mode.*
- *When the firewall serves as an ALG for the Session Initiation Protocol (SIP), by default it performs NAT on the payload and opens dynamic pinholes for media ports. In some cases, depending on the SIP applications in use in your environment, the SIP endpoints have NAT intelligence embedded in their clients. In such cases, you might need to disable the SIP ALG functionality to prevent the firewall from modifying the signaling sessions. When SIP ALG is disabled, if App-ID determines that a session is SIP, the payload is not translated and dynamic pinholes are not opened. See* Disable the SIP Application-level Gateway (ALG).

The following table lists IPv4, NAT, IPv6, NPTv6 and NAT64 ALGs and indicates with a check mark whether the ALG supports each protocol (such as SIP).

App-ID	IPv4	NAT	IPv6	NPTv6	NAT64
SIP	✓	✓	✓	–	–
SCCP	✓	✓	✓	–	–
MGCP	✓	✓	–	–	–
FTP	✓	✓	✓	✓	–
RTSP	✓	✓	✓	✓	–
MySQL	✓	✓	–	–	–
Oracle/SQLNet/ TNS	✓	✓	✓	✓	–
RPC	✓	✓	–	–	–
RSH	✓	✓	–	–	–
UNIStim	✓	✓	–	–	–

© 2018 Palo Alto Networks, Inc.

App-ID	IPv4	NAT	IPv6	NPTv6	NAT64
H.225	✓	✓	—	—	—
H.248	✓	✓	—	—	—

© 2018 Palo Alto Networks, Inc.

Disable the SIP Application-level Gateway (ALG)

The Palo Alto Networks firewall uses the Session Initiation Protocol (SIP) application-level gateway (ALG) to open dynamic pinholes in the firewall where NAT is enabled. However, some applications—such as VoIP—have NAT intelligence embedded in the client application. In these cases, the SIP ALG on the firewall can interfere with the signaling sessions and cause the client application to stop working.

One solution to this problem is to define an Application Override Policy for SIP, but using this approach disables the App-ID and threat detection functionality. A better approach is to disable the SIP ALG, which does not disable App-ID or threat detection.

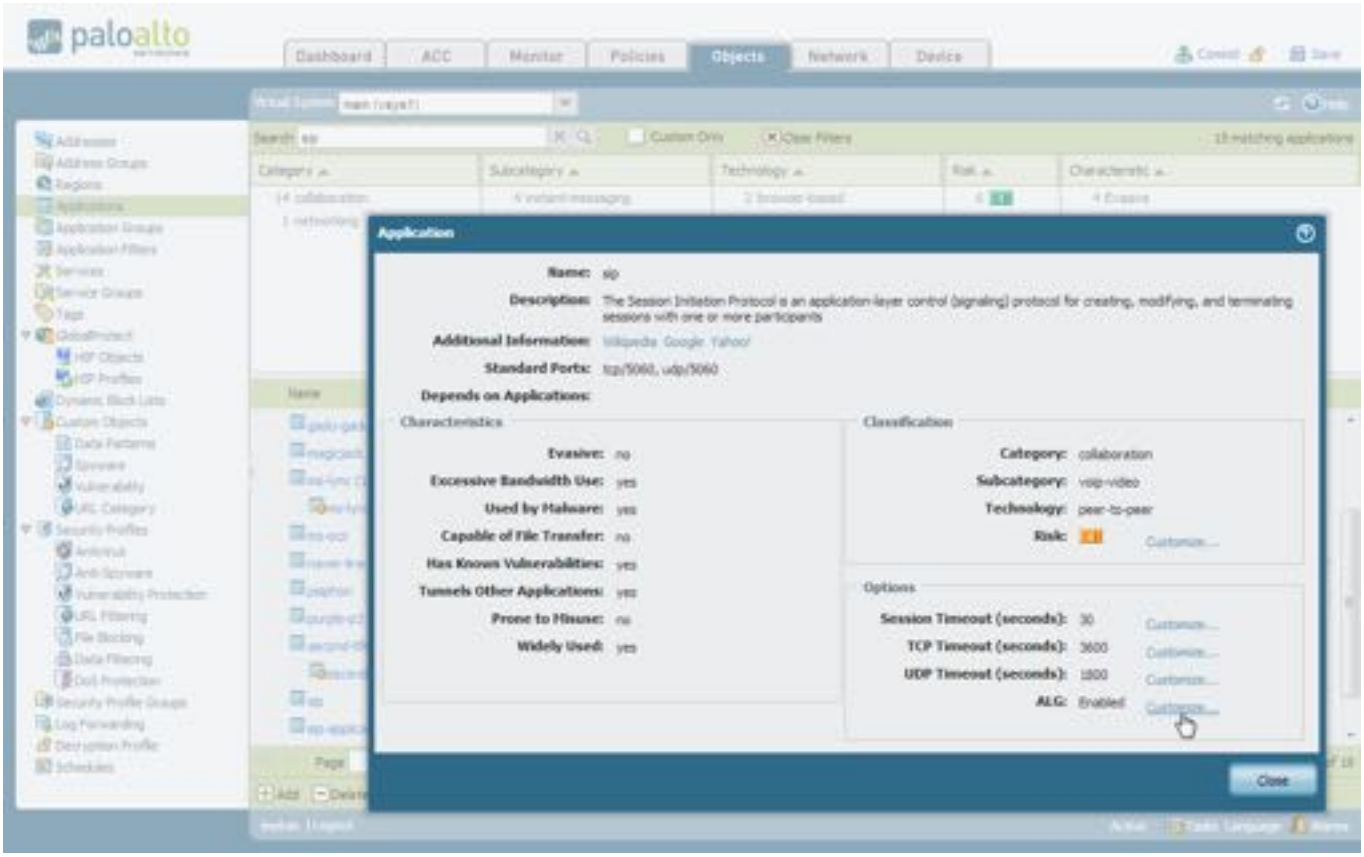
The following procedure describes how to disable the SIP ALG.

STEP 1 | Select **Objects** > **Applications**.

STEP 2 | Select the **sip** application.

You can type **sip** in the **Search** box to help find the sip application.

STEP 3 | Select **Customize...** for **ALG** in the Options section of the Application dialog box.



STEP 4 | Select the **Disable ALG** check box in the Application - sip dialog box and click **OK**.

© 2018 Palo Alto Networks, Inc.



STEP 5 | **Close** the Application dialog box and **Commit** the change.

© 2018 Palo Alto Networks, Inc.

Threat Prevention

The Palo Alto Networks® next-generation firewall protects and defends your network from commodity threats and advanced persistent threats (APTs). The multi-pronged detection mechanisms of the firewall include a signature-based (IPS/Command and Control/Antivirus) approach, heuristics-based (bot detection) approach, sandbox-based (WildFire) approach, and Layer 7 protocol analysis-based (App-ID) approach.

Commodity threats are exploits that are less sophisticated and more easily detected and prevented using a combination of antivirus, anti-spyware, and vulnerability protection features along with URL filtering and Application identification capabilities on the firewall.

Advanced threats are perpetuated by organized cyber adversaries who use sophisticated attack vectors to target your network, most commonly for intellectual property theft and financial data theft. These threats are more evasive and require intelligent monitoring mechanisms for detailed host and network forensics on malware. The Palo Alto Networks next-generation firewall together with WildFire™ and Panorama™ provide a comprehensive solution that intercepts and breaks the attack chain and provides visibility to prevent security infringement on your network infrastructure—both mobile and virtualized.

> Set Up Antivirus, Anti-Spyware, and Vulnerability Protection
> Create Threat Exceptions
> Set Up Data Filtering
> Set Up File Blocking
> Prevent Brute Force Attacks
> Customize the Action and Trigger Conditions for a Brute Force Signature
> Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions
> Enable Evasion Signatures
> Best Practices for Application and Threat Content Updates
> Prevent Credential Phishing
> Share Threat Intelligence with Palo Alto Networks
> Use DNS Queries to Identify Infected Hosts on the Network
> Monitor Blocked IP Addresses
> Learn More About and Assess Threats
> Content Delivery Network Infrastructure for Dynamic Updates
> Threat Prevention Resources

© 2018 Palo Alto Networks, Inc.

Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

Every Palo Alto Networks next-generation firewall comes with predefined Antivirus, Anti-Spyware, and Vulnerability Protection profiles that you can attach to Security policy rules. There is one predefined Antivirus profile, **default**, which uses the default action for each protocol (block HTTP, FTP, and SMB traffic and alert on SMTP, IMAP, and POP3 traffic). There are two predefined Anti-Spyware and Vulnerability Protection profiles:

- **default**—Applies the default action to all client and server critical, high, and medium severity spyware/vulnerability protection events. It does not detect low and informational events.
- **strict**—Applies the block response to all client and server critical, high and medium severity spyware/vulnerability protection events and uses the default action for low and informational events.

To ensure that the traffic entering your network is free from threats, attach the predefined profiles to your basic web access policies. As you monitor the traffic on your network and expand your policy rulebase, you can then design more granular profiles to address your specific security needs.

Use the following workflow to set up the default Antivirus, Anti-Spyware, and Vulnerability Protection Security Profiles.



Palo Alto Networks defines a default action for all anti-spyware and vulnerability protection signatures. To see the default action, select Objects > Security Profiles > Anti-Spyware or Objects > Security Profiles > Vulnerability Protection and then select a profile. Click the Exceptions tab and then click Show all signatures to view the list of the signatures and the corresponding default Action. To change the default action, create a new profile and specify an Action, and/or add individual signature exceptions to Exceptions in the profile.

STEP 1 | Verify that you have a Threat Prevention subscription.

The Threat Prevention subscription bundles the antivirus, anti-spyware, and vulnerability protection features in one license. To verify that you have an active Threat Prevention subscription, select **Device** > **Licenses** and verify that the **Threat Prevention** expiration date is in the future.



STEP 2 | Download the latest content.

1. Select **Device** > **Dynamic Updates** and click **Check Now** at the bottom of the page to retrieve the latest signatures.
2. In the **Actions** column, click **Download** and install the latest Antivirus updates and then download and then **Install** the latest Applications and Threats updates.

STEP 3 | Schedule content updates.



As a best practice, schedule the firewall to download and install Antivirus updates daily and Applications and Threats updates weekly.

© 2018 Palo Alto Networks, Inc.

1. Select **Device** > **Dynamic Updates** and then click **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.
2. Specify the frequency and timing for the updates:
   - **download-only**—The firewall automatically downloads the latest updates per the schedule you define but you must manually **Install** them.
   - **download-and-install**—The firewall automatically downloads and installs the updates per the schedule you define.
3. Click **OK** to save the update schedule; a commit is not required.
4. (Optional) Define a **Threshold** to indicate the minimum number of hours after an update becomes available before the firewall will download it. For example, setting the **Threshold** to **10** means the firewall will not download an update until it is at least 10 hours old regardless of the schedule.
5. (HA only) Decide whether to **Sync To Peer**, which enables peers to synchronize content updates after download and install (the update schedule does not sync across peers; you must manually configure the schedule on both peers).

   There are additional considerations for deciding if and how to **Sync To Peer** depending on your HA deployment:

   - **Active/Passive HA**—If the firewalls are using the MGT port for content updates, then schedule both firewalls to download and install updates independently. However, if the firewalls are using a data port for content updates, then the passive firewall will not download or install updates unless and until it becomes active. To keep the schedules in sync on both firewalls when using a data port for updates, schedule updates on both firewalls and then enable **Sync To Peer** so that whichever firewall is active downloads and installs the updates and also pushes the updates to the passive firewall.
   - **Active/Active HA**—If the firewalls are using the MGT interface for content updates, then select **download-and-install** on both firewalls but do not enable **Sync To Peer**. However, if the firewalls are using a data port, then select **download-and-install** on both firewalls and enable **Sync To Peer** so that if one firewall goes into the active-secondary state, the active-primary firewall will download and install the updates and push them to the active-secondary firewall.

STEP 4 | (Optional) Create custom security profiles for antivirus, anti-spyware, and vulnerability protection.

Alternatively, you can use the predefined default or strict profiles.



Create best practice Security profiles *for the best security posture.*

- To create custom Antivirus Profiles, select **Objects** > **Security Profiles** > **Antivirus** and **Add** a new profile.
- To create custom Anti-Spyware Profiles, select **Objects** > **Security Profiles** > **Anti-Spyware** and **Add** a new profile.
- To create custom Vulnerability Protection Profiles, select **Objects** > **Security Profiles** > **Vulnerability Protection** and **Add** a new profile.

STEP 5 | Attach security profiles to your Security policy rules.

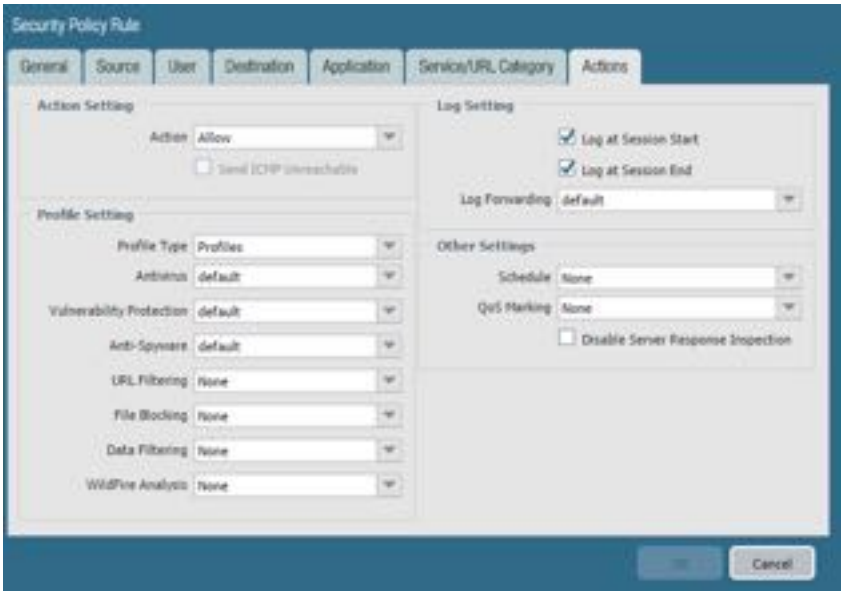


When you configure the firewall with a Security policy rule that uses a Vulnerability Protection profile to block connections, the firewall automatically blocks that traffic in hardware (see Monitor Blocked IP Addresses*).*

1. Select **Policies** > **Security** and select the rule you want to modify.

© 2018 Palo Alto Networks, Inc.

2. In the **Actions** tab, select **Profiles** as the **Profile Type**.
3. Select the security profiles you created for **Antivirus**, **Anti-Spyware**, and **Vulnerability Protection**.



STEP 6 | Commit your changes.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Create Threat Exceptions

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 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:

- Get the latest Antivirus, Threats and Applications, and WildFire signature updates.
- Set Up Antivirus, Anti-Spyware, and Vulnerability Protection and apply these security profiles to your security policy.

STEP 1 | Exclude antivirus signatures from enforcement.



While you can use an Antivirus profile to exclude antivirus signatures from enforcement, you cannot change the action the firewall enforces for a specific antivirus signature. However, you can define the action for the firewall to enforce for viruses found in different types of traffic by editing the Decoders (Objects > Security Profiles > Antivirus > <antivirus-profile> > Antivirus).

1. Select **Objects** > **Security Profiles** > **Antivirus**.
2. **Add** or modify an existing Antivirus profile from which you want to exclude a threat signature and select **Virus Exception**.
3. **Add** the **Threat ID** for the threat signature you want to exclude from enforcement.

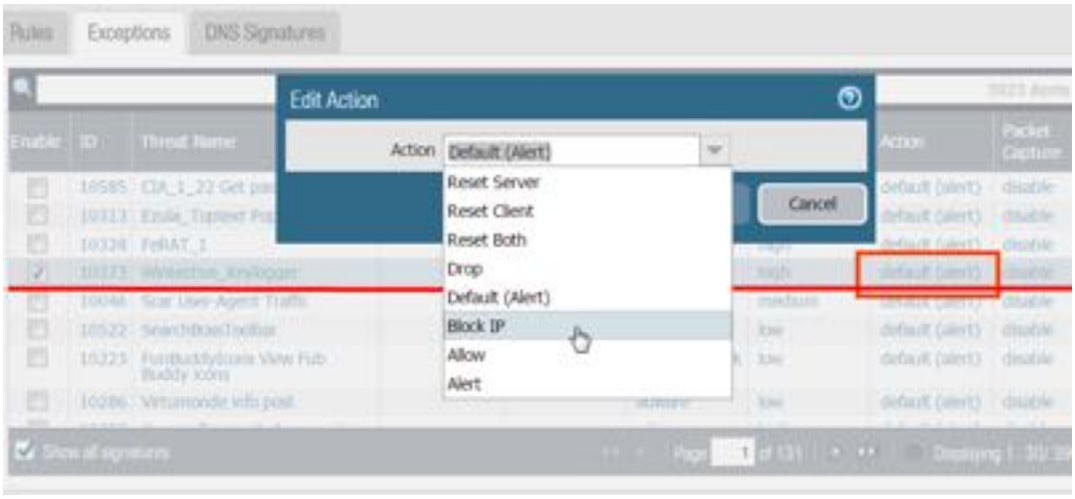


4. Click **OK** to save the Antivirus profile.

STEP 2 | Modify enforcement for vulnerability and spyware signatures (except DNS signatures; skip to the next option to modify enforcement for DNS signatures, which are a type of spyware signature).

1. Select **Objects** > **Security Profiles** > **Anti-Spyware** or **Objects** > **Security Profiles** > **Vulnerability Protection**.
2. **Add** or modify an existing Anti-Spyware or Vulnerability Protection profile from which you want to exclude the threat signature and then select **Exceptions**.
3. **Show all signatures** and then filter to select the signature for which you want to modify enforcement rules.
4. **Enable** the signature for which you want to modify enforcement.
5. Select the **Action** you want the firewall to enforce for this threat signature.

© 2018 Palo Alto Networks, Inc.



For signatures that you want to exclude from enforcement because they trigger false positives, set the **Action** to **Allow**.

6. Click **OK** to save your new or modified Anti-Spyware or Vulnerability Protection profile.

STEP 3 | Modify enforcement for DNS signatures.

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

1. Select **Objects** > **Security Profiles** > **Anti-Spyware**.
2. **Add** or modify the Anti-Spyware profile from which you want to exclude the threat signature, and select **DNS Signatures**.
3. **Add** the **DNS Threat ID** for the DNS signature that you want to exclude from enforcement:



4. Click **OK** to save your new or modified Anti-Spyware profile.

© 2018 Palo Alto Networks, Inc.

Set Up Data Filtering

Use Data Filtering Profiles to prevent sensitive, confidential, and proprietary information from leaving your network. First, create a data pattern to define the information types for which you want the firewall to filter. Predefined patterns and built-in settings make it easy for you to create custom patterns for filtering on social security and credit card numbers or on file properties, such as a document title or author. Continue to add one or more data pattern to a Data Filtering profile and then attach the profile to a Security policy rule to enable data filtering.

If you're using a third-party, endpoint data loss prevention (DLP) solution that populates file properties to indicate sensitive content, then data filtering enables the firewall to enforce your DLP policy. To secure this confidential data, create a custom data pattern to identify the file properties and values tagged by your DLP solution and then log or block the files that your Data Filtering profile detects based on that pattern.

STEP 1 | Define a new data pattern object to detect the information you want to filter.

1. Select **Objects** > **Custom Objects** > **Data Patterns** and **Add** a new object.
2. Provide a descriptive **Name** for the new object.
3. (Optional) Select **Shared** if you want the data pattern to be available to:
   - **Every virtual system (vsys) on a multi-vsys firewall**—If cleared (disabled), the data pattern is available only to the Virtual System selected in the **Objects** tab.
   - **Every device group on Panorama**—If cleared (disabled), the data pattern is available only to the Device Group selected in the **Objects** tab.
4. (Optional—Panorama only) Select **Disable override** to prevent administrators from overriding the settings of this data pattern object in device groups that inherit the object. This selection is cleared by default, which means administrators can override the settings for any device group that inherits the object.
5. (Optional—Panorama only) Select **Data Capture** to automatically collect the data that is blocked by the filter.

   Specify a password for Manage Data Protection on the Settings page to view your captured data (Device > Setup > Content-ID > Manage Data Protection).
6. Set the **Pattern Type** to one of the following:
   - **Predefined**—Filter for credit card and social security numbers.
   - **Regular Expression**—Filter for custom data patterns.
   - **File Properties**—Filter based on file properties and the associated values.
7. **Add** a new rule to the data pattern object.
8. Specify the data pattern according to the **Pattern Type** you selected for this object:
   - **Predefined**—Select the **Name**: either **Credit Card Numbers** or **Social Security Numbers** (with or without dash separator).
   - **Regular Expression**—Specify a descriptive **Name**, select the **File Type** (or types) you want to scan, and then enter the specific **Data Pattern** you want the firewall to detect.
   - **File Properties**—Specify a descriptive **Name**, select the **File Type** and **File Property** you want to scan, and enter the specific **Property Value** that you want the firewall to detect.
9. Click **OK** to save the data pattern.

STEP 2 | Add the data pattern object to a data filtering profile.

1. Select **Objects** > **Security Profiles** > **Data Filtering** and **Add** or modify a data filtering profile.
2. **Add** a new profile rule and select the Data Pattern you created in Step 1.

© 2018 Palo Alto Networks, Inc.

3. Specify **Applications**, **File Types**, and what **Direction** of traffic (upload or download) you want to filter based on the data pattern.

   The file type you select must be the same file type you defined for the data pattern in Step 1 or it must be a file type that includes the data pattern file type. For example, you could define both the data pattern object and the data filtering profile to scan all Microsoft Office documents. Or, you could define the data pattern object to match to only Microsoft PowerPoint Presentations while the data filtering profile scan all Microsoft Office documents.

   If a data pattern object is attached to a data filtering profile and the configured file types do not align between the two, the profile will not correctly filter documents matched to the data pattern object.
4. Set the **Alert Threshold** to specify the number of times the data pattern must be detected in a file to trigger an alert.
5. Set the **Block Threshold** to block files that contain at least this many instances of the data pattern.
6. Set the **Log Severity** recorded for files that match this rule.
7. Click **OK** to save the data filtering profile.

STEP 3 | Apply the data filtering settings to traffic.

1. Select **Policies** > **Security** and **Add** or modify a security policy rule.
2. Select **Actions** and set the Profile Type to **Profiles**.
3. Attach the Data Filtering profile you created in Step 2 to the security policy rule.
4. Click **OK**.

STEP 4 | (Recommended) Prevent web browsers from resuming sessions that the firewall has terminated.

This option ensures that when the firewall detects and then drops a sensitive file, a web browser cannot resume the session in an attempt to retrieve the file.

1. Select **Device** > **Setup** > **Content-ID** and edit Content-ID Settings.
2. Clear the **Allow HTTP header range option**.
3. Click **OK**.

STEP 5 | Monitor files that the firewall is filtering.

Select **Monitor** > **Data Filtering** to view the files that the firewall has detected and blocked based on your data filtering settings.

© 2018 Palo Alto Networks, Inc.

Set Up File Blocking

File Blocking Profiles allow you to identify specific file types that you want to want to block or monitor. For most traffic (including traffic on your internal network) you will want to block files that are known to carry threats or that have no real use case for upload/download. Currently, these include batch files, DLLs, Java class files, help files, Windows shortcuts (.lnk), and BitTorrent files. Additionally, to provide drive-by download protection, allow download/upload of executables and archive files (.zip and .rar), but force users to acknowledge that they are transferring a file so that they will notice that the browser is attempting to download something they were not aware of. For policy rules that allow general web browsing, be more strict with your file blocking because the risk of users unknowingly downloading malicious files is much higher. For this type of traffic you will want to attach a more strict file blocking profile that also blocks portable executable (PE) files.

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 quickly enable best practice file blocking settings:

- **basic file blocking**—Attach this profile to the Security policy rules that allow traffic to and from less sensitive applications to block files that are commonly included in malware attack campaigns or that have no real use case for upload/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, .bat. Additionally, it prompts users to acknowledge when they attempt to download encrypted-rar or encrypted-zip files. This rule alerts on all other file types to give you complete visibility into all file types coming in and out of your network.
- **strict file blocking**—Use this stricter profile on the Security policy rules that allow access to your most sensitive applications. This profile blocks the same file types as the other profile, and additionally blocks flash, .tar, multi-level encoding, .cab, .msi, encrypted-rar, and encrypted-zip files.

These predefined profiles are designed to provide the most secure posture for your network. However, if you have business-critical applications that rely on some of the applications that are blocked in these default profiles, you can clone the profiles and modify them as necessary. Make sure that you only use the modified profiles for those users who need to upload and/or download a risky file type. Additionally, to reduce your attack surface, make sure you are using other security measures to ensure that the files your users are uploading and downloading do not pose a threat to your organization. For example, if you must allow download of PE files, make sure you are sending all unknown PE files to WildFire foranalysis. Additionally, maintain a strict URL filtering policy to ensure that users cannot download content from web sites that have been known to host malicious content.

STEP 1 | Create the file blocking profile.

1. Select **Objects** > **Security Profiles** > **File Blocking** and **Add** a profile.
2. Enter a **Name** for the file blocking profile such as **`Block_EXE`**.
3. (Optional) Enter a **Description**, such as **`Block users from downloading exe files from websites`**.
4. (Optional) Specify that the profile is **Shared** with:
   - **Every virtual system (vsys) on a multi-vsys firewall**—If cleared (disabled), the profile is available only to the Virtual System selected in the **Objects** tab.
   - **Every device group on Panorama**—If cleared (disabled), the profile is available only to the Device Group selected in the **Objects** tab.
5. (Optional—Panorama only) Select **Disable override** to prevent administrators from overriding the settings of this file blocking profile in device groups that inherit the profile. This selection is cleared

© 2018 Palo Alto Networks, Inc.

by default, which means administrators can override the settings for any device group that inherits the profile.

STEP 2 | Configure the file blocking options.

1. **Add** and define a rule for the profile.
2. Enter a **Name** for the rule, such as **`BlockEXE`**.
3. Select **Any** or specify one or more specific **Applications** for filtering, such as **web-browsing**.

   Only web browsers can display the response page (continue prompt) that allows users to confirm their Choosing any other application results in blocked traffic for those applications because there is no prompt displayed to allow users to continue.
4. Select **Any** or specify one or more specific **File Types**, such as **exe**.
5. Specify the **Direction**, such as **download**.
6. Specify the **Action** (**alert**, **block**, or **continue**). For example, select **continue** to prompt users for confirmation before they are allowed to download an executable (.exe) file. Alternatively, you could **block** the specified files or you could configure the firewall to simply trigger an **alert** when a user downloads an executable file.
7. Click **OK** to save the profile.

STEP 3 | Apply the file blocking profile to a security policy rule.

1. Select **Policies** > **Security** and either select an existing policy rule or **Add** a new rule as described in Set Up a Basic Security Policy.
2. On the **Actions** tab, select the file blocking profile you configured in the previous step. In this example, the profile name is **`Block_EXE`**.
3. **Commit** your configuration.

STEP 4 | To test your file blocking configuration, access an endpoint PC in the trust zone of the firewall and attempt to download an executable file from a website in the untrust zone; a response page should display. Click **Continue** to confirm that you can download the file. You can also set other actions, such as **alert** or **block**, which do not provide an option for the user to continue the download. The following shows the default response page for File Blocking:



STEP 5 | (Optional) Define custom file blocking response pages (**Device** > **Response Pages**). This allows you to provide more information to users when they see a response page. You can include information such as company policy information and contact information for a Helpdesk.

When you create a file blocking profile with the continue action, you can choose only the web-browsing application. If you choose any other application, traffic that matches the security policy will not flow through the firewall because users are not prompted with an option to continue. Additionally, you need to configure and enable a decryption policy for HTTPS websites.

© 2018 Palo Alto Networks, Inc.



Check your logs to determine the application used when you test this feature. For example, if you are using Microsoft SharePoint to download files, even though you are using a web-browser to access the site, the application is actually `sharepoint-base`*, or* `sharepoint-document`*. (It can help to set the application type to Any for testing.)*

© 2018 Palo Alto Networks, Inc.

Prevent Brute Force Attacks

A brute force attack uses a large volume of requests/responses from the same source or destination IP address to break into a system. The attacker employs a trial-and-error method to guess the response to a challenge or a request.

The Vulnerability Protection profile on the firewall includes signatures to protect you from brute force attacks. Each signature has an ID, Threat Name, and Severity and is triggered when a pattern is recorded. The pattern specifies the conditions and interval at which the traffic is identified as a brute-force attack; some signatures are associated with another child signature that is of a lower severity and specifies the pattern to match against. When a pattern matches against the signature or child signature, it triggers the default action for the signature.

To enforce protection:

- Attach the Vulnerability Protection profile to a Security policy rule. See Set Up Antivirus, Anti-Spyware, and Vulnerability Protection.
- Install content updates that include new signatures to protect against emerging threats. See Install Content and Software Updates.

© 2018 Palo Alto Networks, Inc.

Customize the Action and Trigger Conditions for a Brute Force Signature

The firewall includes two types of predefined brute force signatures—parent signatures and child signatures. A child signature is a single occurrence of a traffic pattern that matches the signature. A parent signature is associated with a child signature and is triggered when multiple events occur within a specified time interval and that matches the traffic pattern defined in the child signature.

Typically, the default action for a child signature is *allow* because a single event is not indicative of an attack. This ensures that legitimate traffic is not blocked and avoids generating threat logs for non-noteworthy events. Palo Alto Networks recommends that you do not change the default action without careful consideration.

In most cases, the brute force signature is a noteworthy event due to its recurrent pattern. If needed, you can do one of the following to customize the action for a brute-force signature:

- Create a rule to modify the default action for all signatures in the brute force category. You can choose to allow, alert, block, reset, or drop the traffic.
- Define an exception for a specific signature. For example, you can search for and define an exception for a CVE.

For a parent signature, you can modify both the trigger conditions and the action; for a child signature, you can modify only the action.



To effectively mitigate an attack, specify the block-ip address action instead of the drop or reset action for most brute force signatures.

STEP 1 | Create a new Vulnerability Protection profile.

1. Select **Objects** > **Security Profiles** > **Vulnerability Protection** and **Add** a profile.
2. Enter a **Name** for the Vulnerability Protection profile.
3. (Optional) Enter a **Description**.
4. (Optional) Specify that the profile is **Shared** with:
   - **Every virtual system (vsys) on a multi-vsys firewall**—If cleared (disabled), the profile is available only to the Virtual System selected in the **Objects** tab.
   - **Every device group on Panorama**—If cleared (disabled), the profile is available only to the Device Group selected in the **Objects** tab.
5. (Optional—Panorama only) Select **Disable override** to prevent administrators from overriding the settings of this Vulnerability Protection profile in device groups that inherit the profile. This selection is cleared by default, which means administrators can override the settings for any device group that inherits the profile.

STEP 2 | Create a rule that defines the action for all signatures in a category.

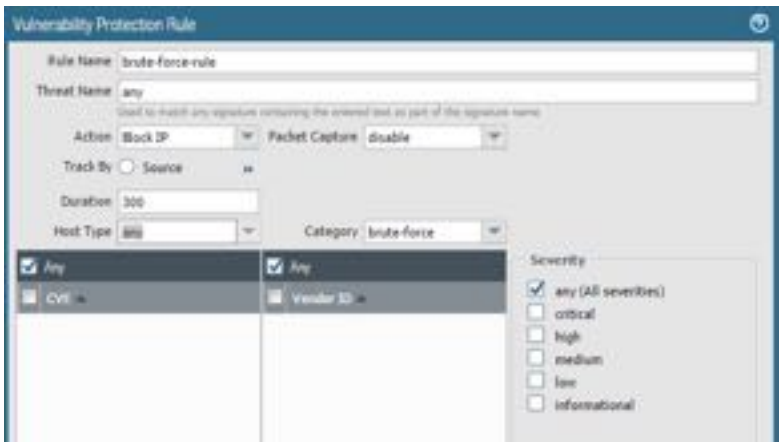
1. On the **Rules** tab, **Add** and enter a **Rule Name** for a new rule.
2. (Optional) Specify a specific threat name (default is **any**).
3. Set the **Action**. In this example, it is set to **Block IP**.



If you set a Vulnerability Protection profile to Block IP, the firewall first uses hardware to block IP addresses. If attack traffic exceeds the blocking capacity of the hardware, the firewall then uses software blocking mechanisms to block the remaining IP addresses.

© 2018 Palo Alto Networks, Inc.

4. Set **Category** to **brute-force**.
5. (Optional) If blocking, specify the **Host Type** on which to block: **server** or **client** (default is **any**).
6. See Step 3 to customize the action for a specific signature.
7. See Step 4 to customize the trigger threshold for a parent signature.

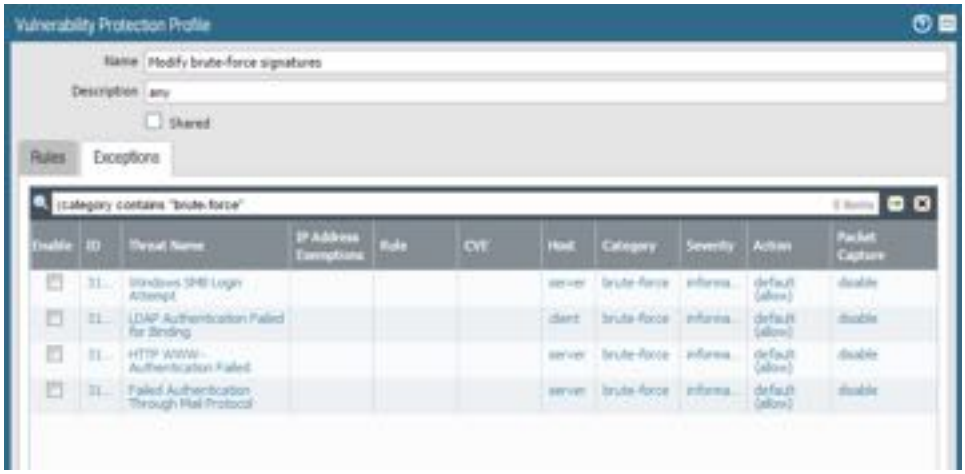


8. Click **OK** to save the rule and the profile.

STEP 3 | (Optional) Customize the action for a specific signature.

1. On the **Exceptions** tab, **Show all signatures** to find the signature you want to modify.

   To view all the signatures in the brute-force category, search for `category contains 'brute-force'`.
2. To edit a specific signature, click the predefined default action in the Action column.



3. Set the action: **Allow**, **Alert**, **Block Ip**, or **Drop**. If you select **Block Ip**, complete these additional tasks:
   1. Specify the **Time** period (in seconds) after which to trigger the action.
   2. Specify whether to **Track By** and block the IP address using the **IP source** or the **IP source and destination**.
4. Click **OK**.
5. For each modified signature, select the check box in the **Enable** column.
6. Click **OK**.

STEP 4 | Customize the trigger conditions for a parent signature.

A parent signature that can be edited is marked with this icon: .

In this example, the search criteria was brute force category and CVE-2008-1447.

© 2018 Palo Alto Networks, Inc.

1. Edit ( ) the time attribute and the aggregation criteria for the signature.
2. To modify the trigger threshold, specify the **Number of Hits** per number of **seconds**.
3. Specify whether to aggregate the number of hits (**Aggregation Criteria**) by **source**, **destination**, or **source-and-destination**.
4. Click **OK**.

STEP 5 | Attach this new profile to a Security policy rule.

1. Select **Policies** > **Security** and **Add** or modify a Security policy rule.
2. On the **Actions** tab, select **Profiles** as the **Profile Type** for the Profile Setting.
3. Select your **Vulnerability Protection** profile.
4. Click **OK**.

STEP 6 | Commit your changes.

1. Click **Commit**.

© 2018 Palo Alto Networks, Inc.

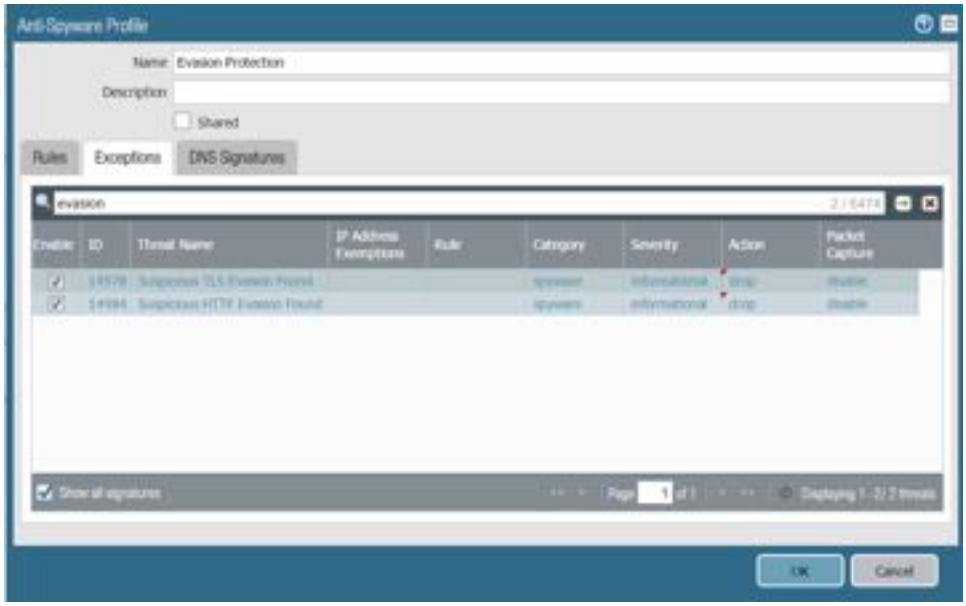
Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions

To monitor and protect your network from most Layer 4 and Layer 7 attacks, here are a few recommendations.

- Upgrade to the most current PAN-OS software version and content release version to ensure that you have the latest security updates. See Install Content and Software Updates.
- Set up the firewall to act as a DNS proxy and enable evasion signatures:
  - Configure a DNS Proxy Object.

    When acting as a DNS proxy, the firewall resolves DNS requests and caches hostname-to-IP address mappings to quickly and efficiently resolve future DNS queries.
  - Enable Evasion Signatures

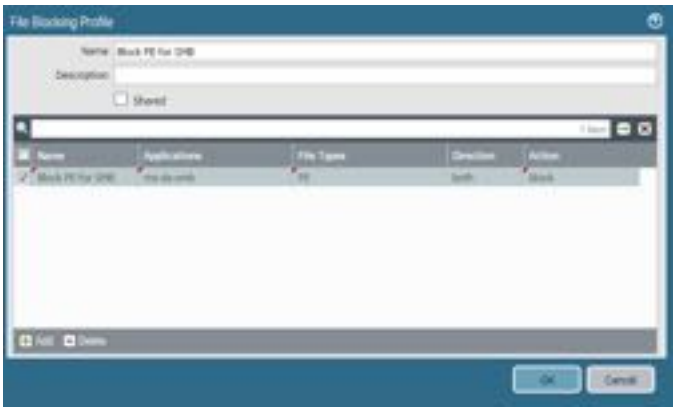
    Evasion signatures that detect crafted HTTP or TLS requests can send alerts when clients connect to a domain other than the domain specified in the original DNS request. Make sure to configure DNS proxy before you enable evasion signatures. Without DNS proxy, evasion signatures can trigger alerts when a DNS server in the DNS load balancing configuration returns different IP addresses—for servers hosting identical resources—to the firewall and client in response to the same DNS request.



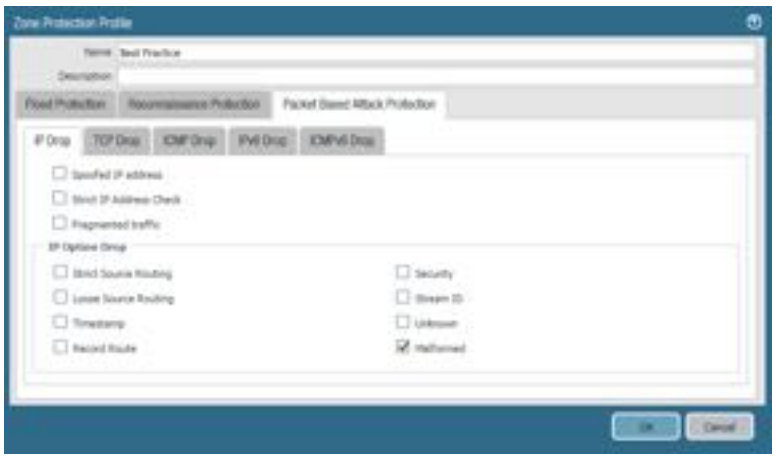
- For servers, create Security policy rules to allow only the application(s) that you sanction on each server. Verify that the standard port for the application matches the listening port on the server. For example, to ensure that only SMTP traffic is allowed to your email server, set the Application to **smtp** and set the Service to **application-default**. If your server uses only a subset of the standard ports (for example, if your SMTP server uses only port 587 while the SMTP application has standard ports defined as 25 and 587), create a new custom service that includes only port 587 and use that new service in your security policy rule instead of application-default. Additionally, make sure you restrict access to specific source and destinations zones and sets of IP addresses.
- Block all unknown applications and traffic using the Security policy. Typically, the only applications classified as unknown traffic are internal or custom applications on your network and potential threats. Unknown traffic can be either non-compliant applications or protocols that are anomalous or abnormal or it can be known applications that are using non-standard ports, both of which should be blocked. See Manage Custom or Unknown Applications.

© 2018 Palo Alto Networks, Inc.

- Set Up File Blocking to block Portable Executable (PE) file types for internet-based SMB (Server Message Block) traffic from traversing trust to untrust zones (ms-ds-smb applications).



- Create a Zone Protection profile that is configured to protect against packet-based attacks (**Network** > **Network Profiles** > **Zone Protection**):
  - Select the option to drop **Malformed** IP packets (**Packet Based Attack Protection** > **IP Drop**).



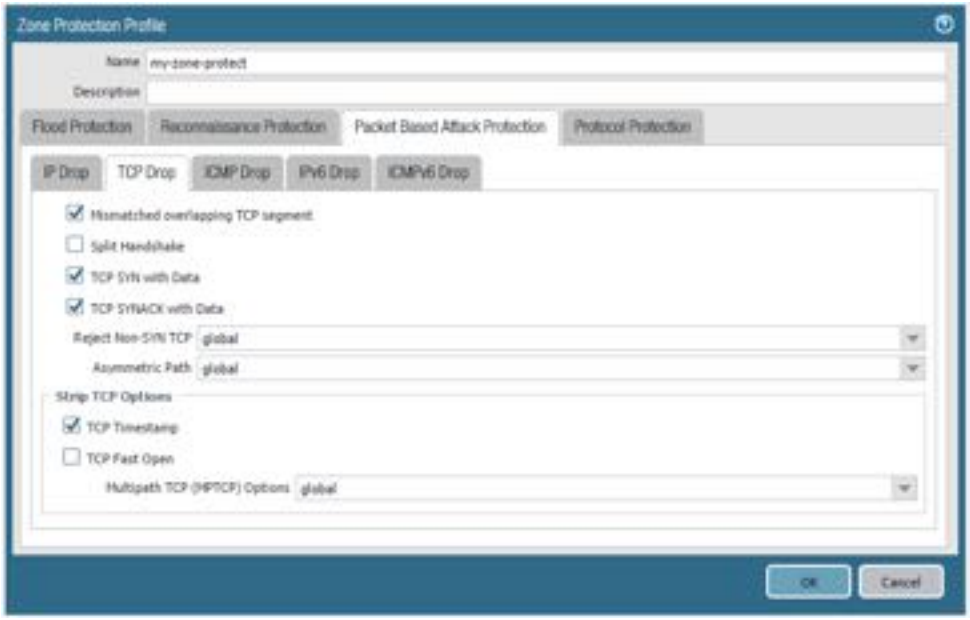
  - Enable the drop **Mismatched overlapping TCP segment** option (**Packet Based Attack Protection** > **TCP Drop**).

    By deliberately constructing connections with overlapping but different data in them, attackers attempt to cause misinterpretation of the intent of the connection and deliberately induce false positives or false negatives. Attackers also use IP spoofing and sequence number prediction to intercept a user's connection and inject their own data into that connection. Selecting the **Mismatched overlapping TCP segment** option specifies that PAN-OS discards frames with mismatched and overlapping data. Received segments are discarded when they are contained within another segment, when they overlap with part of another segment, or when they contain another complete segment.
  - Enable the drop **TCP SYN with Data** and drop **TCP SYNACK with Data** options (**Packet Based Attack Protection** > **TCP Drop**).

    Dropping SYN and SYN-ACK packets that contain data in the payload during a three-way handshake increases security by blocking malware contained in the payload and preventing it from extracting unauthorized data before the TCP handshake is completed.
  - Strip TCP timestamps from SYN packets before the firewall forwards the packet (**Packet Based Attack Protection** > **TCP Drop**).

© 2018 Palo Alto Networks, Inc.

When you enable the **Strip TCP Options - TCP Timestamp** option, the TCP stack on both ends of the TCP connection will not support TCP timestamps. This prevents attacks that use different timestamps on multiple packets for the same sequence number.



- If you configure IPv6 addresses on your network hosts, be sure to enable support for IPv6 if not already enabled (**Network** > **Interfaces** > **Ethernet** > **IPv6**).

  Enabling support for IPv6 allows access to IPv6 hosts and also filters IPv6 packets encapsulated in IPv4 packets, which prevents IPv6 over IPv4 multicast addresses from being leveraged for network reconnaissance.



- Enable support for multicast traffic so that the firewall can enforce policy on multicast traffic (**Network** > **Virtual Router** > **Multicast**).



- Disable the options to **Forward datagrams exceeding UDP content inspection queue** and **Forward segments exceeding TCP content inspection queue** (**Device** > **Setup** > **Content-ID** > **Content-ID Settings**).

  By default, when the TCP or UDP content inspection queues are full, the firewall skips content inspection for TCP segments or UDP datagrams that exceed the queue limit of 64. Disabling this option ensures content inspection for all TCP and UDP datagrams that the firewall allows. Only under very

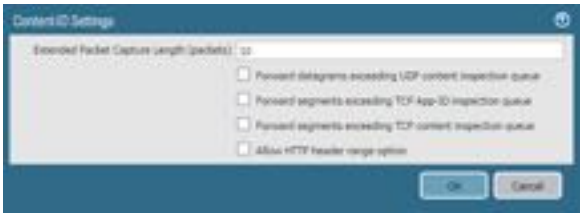
© 2018 Palo Alto Networks, Inc.

specific circumstances—for example, if the firewall platform is not sized appropriately to align with a use case—could disabling these settings impact performance.

- Disable the **Allow HTTP header range option** (**Device** > **Setup** > **Content-ID** > **Content-ID Settings**).

  The HTTP header range option allows a client to fetch only part of a file. When a next-generation firewall in the path of a transfer identifies and drops a malicious file, it terminates the TCP session with an RST packet. If the web browser implements the HTTP header range option, it can start a new session to fetch only the remaining part of the file, which prevents the firewall from triggering the same signature again due to the lack of context into the initial session and, at the same time, allows the web browser to reassemble the file and deliver the malicious content. Disabling this option prevents this from happening.

  Disabling this option should not impact device performance. However, HTTP file transfer interruption recovery may be impaired. In addition, disabling this option can impact streaming media services, such as Netflix, Windows Server Updates Services (WSUS), and Palo Alto Networks content updates.



- Create a Vulnerability Protection Profile that blocks protocol anomalies and blocks vulnerabilities with low and high severities.

  A protocol anomaly occurs when a protocol behavior deviates from standard and compliant usage. For example, a malformed packet, poorly written application, or an application running on a non-standard port would all be considered protocol anomalies, and could be used as evasion tools.

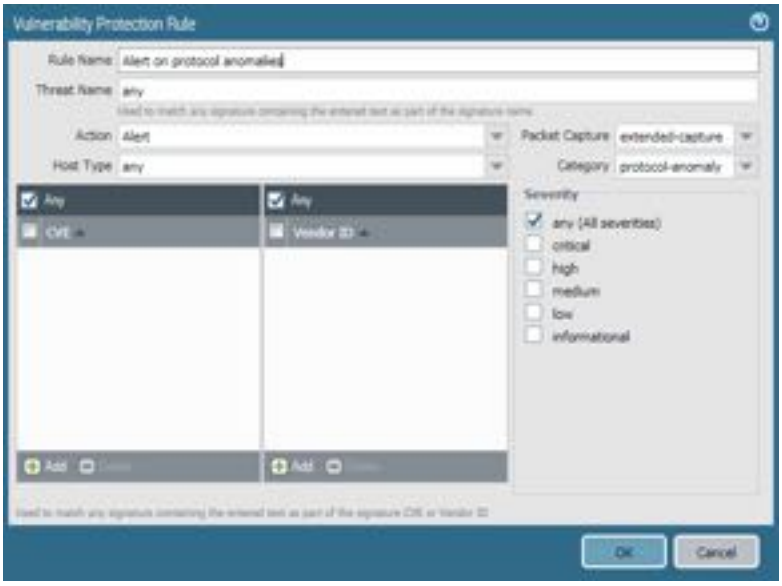
  If yours is a mission-critical network, where the business's highest priority is application availability, you should begin by alerting on protocol anomalies for a period of time to ensure that no critical internal applications are using established protocols in a non-standard way. If you find that certain critical applications trigger protocol anomaly signatures, you can then exclude those applications from protocol anomaly enforcement. To do this, add another rule to the Vulnerability Protection Profile that whitelists protocol anomalies and attach the profile to the security policy rule that enforces traffic to and from the critical applications.

  Make sure that Vulnerability Protection Profile rules and security policy rules that whitelist protocol anomalies for critical internal applications are listed above rules that block protocol anomalies. Traffic is evaluated against security policy rules and associated Vulnerability Protection Profiles rules from top to bottom, and is enforced based on the first matching rule.

  - Begin by alerting on protocol anomalies:

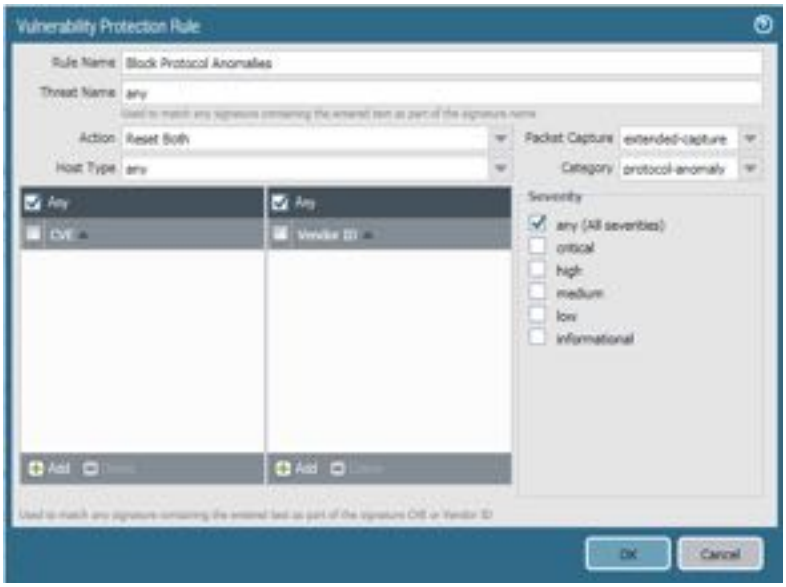
    Create a Vulnerability Protection Profile rule with the **Action** set to Alert, the **Category** set to protocol-anomaly, and the **Severity** set to Any. Monitor your traffic to determine if any critical internal applications are using established protocols in non-standard ways. If you find this to be true, continue to whitelist protocol anomalies for those applications, and then block protocol anomalies for all other applications.

© 2018 Palo Alto Networks, Inc.



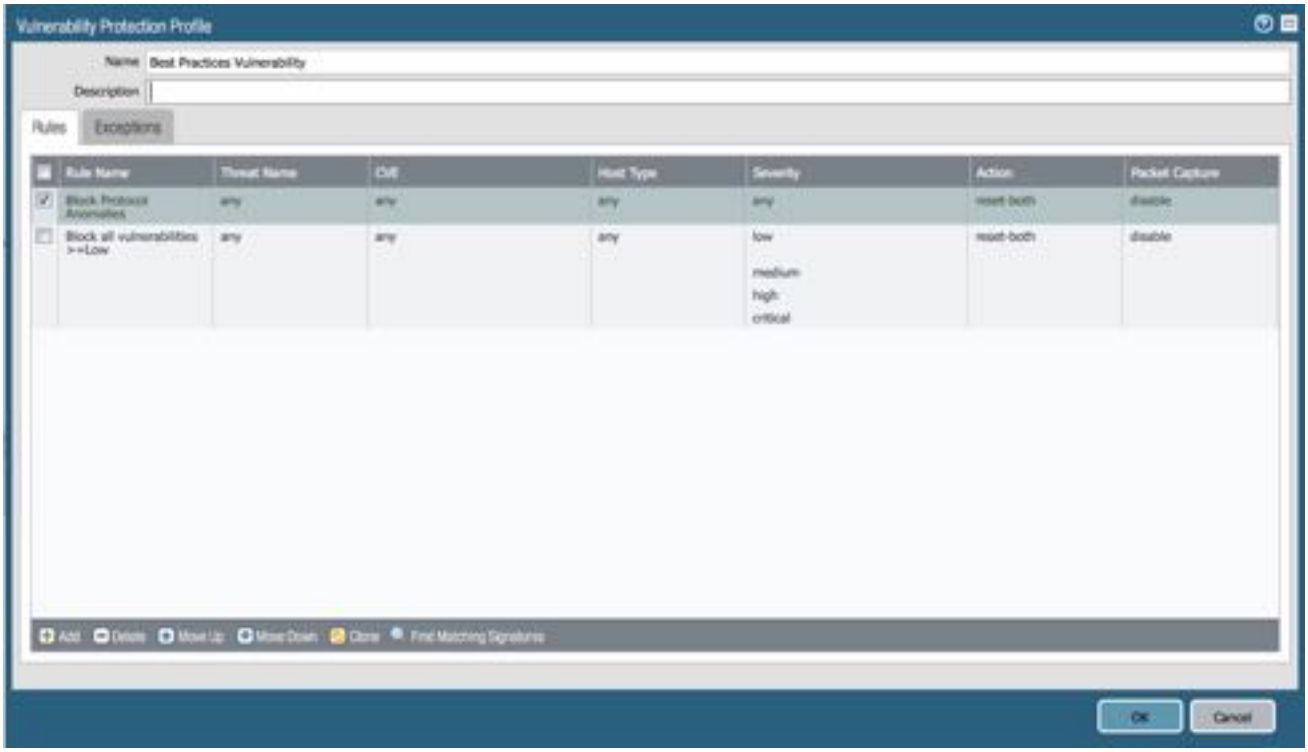
- Block protocol anomalies:

  Create a Vulnerability Protection Profile rule with the **Category** set to protocol-anomaly, the rule **Action** set to Reset Both, and the **Severity** set to Any.



- Optionally whitelist protocol anomalies for critical applications that use established protocols in a non-standard way. To do this, create a Vulnerability Protection Profile rule that allows protocol anomalies: set the rule **Action** to Allow, the **Category** to protocol-anomaly, and the **Severity** to any. Attach the Vulnerability Protection Profile rule to the security policy rule that enforces traffic to and from critical applications.
- Add another rule to the Vulnerability Protection Profile to block all vulnerabilities with low and higher severity. This rule must be listed after the rule that blocks protocol anomalies.

© 2018 Palo Alto Networks, Inc.



- Continue to attach the following security profiles to your Security policy rules to provide signature-based protection:
  - An Anti-Spyware profile to block all spyware with severity low and higher.
  - An Antivirus profile to block all content that matches an antivirus signature.

© 2018 Palo Alto Networks, Inc.

Enable Evasion Signatures

Palo Alto Networks evasion signatures detect crafted HTTP or TLS requests, and can alert to instances where a client connects to a domain other than the domain specified in a DNS query. Evasion signatures are effective only when the firewall is also enabled to act as a DNS proxy and resolve domain name queries. As a best practice, take the following steps to enable evasion signatures.

STEP 1 | Enable a firewall intermediate to clients and servers to act as a DNS proxy.

Configurea DNS Proxy Object, including:

- Specify the interfaces on which you want the firewall to listen for DNS queries.
- Define the DNS servers with which the firewall communicates to resolve DNS requests.
- Set up static FQDN-to-IP address entries that the firewall can resolve locally, without reaching out to DNS servers.
- Enable caching for resolved hostname-to-IP-address mappings.

STEP 2 | Get the latest Applications and Threats content version (at least content version 579 or later).

1. Select **Device** > **Dynamic Updates**.
2. **Check Now** to get the latest Applications and Threats content update.
3. Download and Install Applications and Threats content version 579 (or later).

STEP 3 | Define how the firewall should enforce traffic matched to evasion signatures.

1. Select **Objects** > **Security Profiles** > **Anti-Spyware** and **Add** or modify an Anti-spyware profile.
2. Select **Exceptions** and select **Show all signatures**.
3. Filter signatures based on the keyword `evasion`.
4. For all evasion signatures, set the **Action** to any setting other than allow or the default action (the default action is for evasion signatures is allow). For example, set the **Action** for signature IDs 14978 and 14984 to **alert** or **drop**.
5. Click **OK** to save the updated Anti-spyware profile.
6. Attach the Anti-spyware profile to a security policy rule: Select **Policies** > **Security**, select the desired policy to modify and then click the **Actions** tab. In Profile Settings, click the drop-down next to **Anti-Spyware** and select the anti-spyware profile you just modified to enforce evasion signatures.

STEP 4 | Commit your changes.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Best Practices for Applications and Threats Content Updates

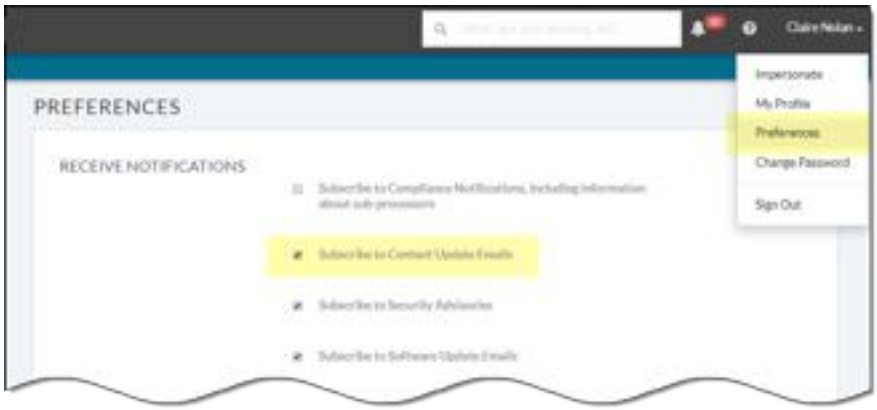
The best practices to deploy content updates help to ensure seamless policy enforcement as new threat signatures and applications are introduced or modified in a content release. Because of the policy impact of new application and threat signatures, consider your network security and availability requirements as you apply best practices:

- An organization with a *security-first* posture prioritizes protection using the latest threat signatures over application availability. You're primarily using the firewall for its threat prevention capabilities.
- A *mission-critical* network prioritizes application availability over protection using the latest threat signatures. Your network has zero tolerance for downtime. The firewall is deployed inline to enforce security policy and if you're using App-ID in security policy, any change to content that affects App-ID could cause downtime.

You can take a mission-critical or security-first approach to deploying content updates, or you can apply a mix of both approaches to meet the needs of the business. Follow these best practices to most effectively absorb the new application and threat signatures that are delivered to the firewall in content updates:

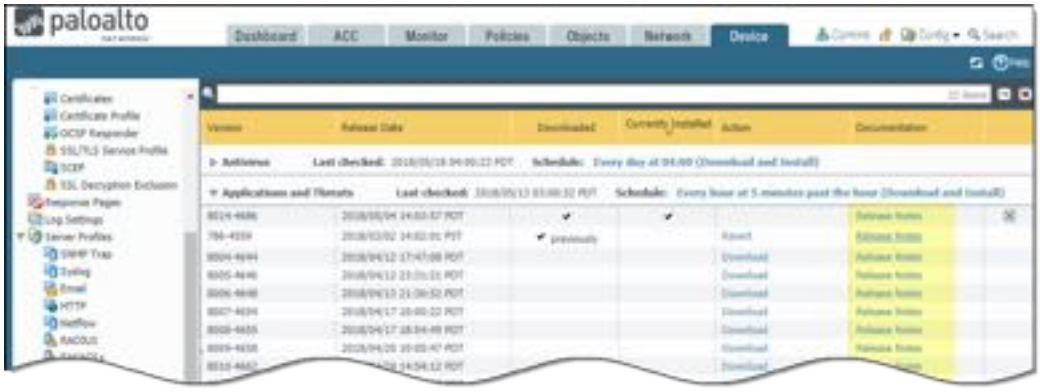
- Always review Content Release Notes for the list of the newly-identified and modified application and threat signatures that the content release introduces. Content Release Notes also describe how the update might impact existing security policy enforcement and provides recommendations on how you can modify your security policy to best leverage what's new.

  To subscribe to get notifications for new content updates, visit the Customer Support Portal, edit your **Preferences**, and select **Subscribe to Content Update Emails**.



  You can also review Content Release Notes for apps and threats on the Palo Alto Networks Support Portal or directly in the firewall web interface: select **Device** > **Dynamic Updates** and open the **Release Note** for a specific content version.

© 2018 Palo Alto Networks, Inc.



The Notes section of Content Release Notes highlights future updates that Palo Alto Networks has identified as possibly significantly impacting coverage: for example, new App-IDs or decoders. Check for these future updates, so that you can account for any policy impact in advance of the release.

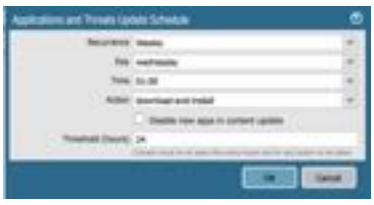
- Schedule content updates so that they download and install automatically and, based on your network security and availability requirements, set a threshold that determines the amount of time the firewall waits before installing the latest content:

  - If you have a security-first posture, do not set a threshold to delay receiving the latest content update. Enable the firewall to download and install content updates as they are made available so that you are always equipped with the most up-to-date threat prevention signatures.
  - If your network is mission-critical, schedule a 24-hour threshold for content updates. This 24-hour delay ensures that the firewall only installs content releases after they have been available and functioning correctly in customer environments for at least 24 hours.

    To mitigate any risk associated with enabling new application and threat signatures, you can stagger the roll-out of new content. Provide the new content to locations with less business risk (fewer users in satellite offices) before deploying them to locations with more business risk (such as locations with critical applications). Confining the latest content updates to certain firewalls before deploying them across your network makes it easier to troubleshoot any issues that arise.

    Use Panorama to push staggered schedules to different firewalls or device groups.

    To schedule content updates, select **Device** > **Dynamic Updates**. Configure the **Schedule** for Applications and Threats content updates, set the schedule **Action** to **download-and-install**, and set (optionally) the **Threshold** to **24** hours.



- Manage New App-IDs Introduced in Content Releases. Always review the new App-IDs that a content release introduces and assess the policy impact of the newly-identified applications. In mission-critical environments, you can wait to install new applications until after reviewing their policy impact. If you cannot modify security policy before installing the latest content update, you can disable new applications in the content update and review policy impact of these applications later.
- If yours is a mission-critical environment, test new Applications and Threats content updates in a dedicated staging environment before enabling them in your production environment. The easiest way

© 2018 Palo Alto Networks, Inc.

to test new applications and threats is to use a test firewall to tap into production traffic. Install the latest content on the test firewall and monitor the firewall as it processes the traffic copied from your production environment. You can also use test clients and a test firewall or packet captures (PCAPs) to simulate production traffic. Using PCAPs works well to simulate traffic for diverse deployments where firewall security policy varies depending on location.

© 2018 Palo Alto Networks, Inc.

Prevent Credential Phishing

Phishing sites are sites that attackers disguise as legitimate websites with the aim to steal user information, especially the credentials that provide access to your network. When a phishing email enters a network, it takes just a single user to click the link and enter credentials to set a breach into motion. You can detect and prevent in-progress phishing attacks, thereby preventing credential theft, by controlling sites to which users can submit corporate credentials based on the site's URL category. This allows you to block users from submitting credentials to untrusted sites while allowing users to continue to submit credentials to corporate and sanctioned sites.

Credential phishing prevention works by scanning username and password submissions to websites and comparing those submissions against valid corporate credentials. You can choose what websites you want to either allow or block corporate credential submissions to based on the URL category of the website. When the firewall detects a user attempting to submit credentials 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 sites classified in certain URL categories, but still allows them 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.

To enable Credential phishing prevention you must configure both User-ID to detect when users submit valid corporate credentials to a site (as opposed to personal credentials) and URL Filtering to specify the URL categories in which you want to prevent users from entering their corporate credentials. The following topics describe the different methods you can use to detect credential submissions and provide instructions for configuring credential phishing protection.

- Methodsto Check for Corporate Credential Submissions
- ConfigureCredential Detection with the Windows-based User-ID Agent
- SetUp Credential Phishing Prevention

Methods to Check for Corporate Credential Submissions

Before you Set Up Credential Phishing Prevention, decide which method you want the firewall to use to check if credentials submitted to a web page are valid, corporate credentials.

Method to Check Submitted Credentials	User-ID Configuration Requirements	How does this method detect corporate usernames and/or passwords as users submit them to websites?
Group Mapping	Group Mapping configuration on the firewall	The firewall determines if the username a user submits to a restricted site matches any valid corporate username. To do this, the firewall matches the submitted username to the list of usernames in its user-to-group mapping table to detect when users submit a corporate usernames to a site in a restricted category. This method only checks for corporate username submissions based on LDAP group membership, which makes it simple to configure, but more prone to false positives.
IP User Mapping	IP-address-to- username mappings	The firewall determines if the username a user submits to a restricted site maps to the IP address of the logged-in user.

© 2018 Palo Alto Networks, Inc.

Method to Check Submitted Credentials	User-ID Configuration Requirements	How does this method detect corporate usernames and/or passwords as users submit them to websites?
	identified through User Mapping, GlobalProtect, or Authentication Policy and Captive Portal.	To do this, the firewall matches the IP address of the logged in user and the username submitted to a web site to its IP-address-to-user mapping table to detect when users submit their corporate usernames to a site in a restricted category. Because this method matches the IP address of the logged-in user associated with the session against the IP-address-to-username mapping table, it is an effective method for detecting corporate username submissions, but it does not detect corporate password submission. If you want to detect corporate username and password submission, you must use the Domain Credential Filter method.
Domain Credential Filter	Windows-based User-ID agent configured with the User-ID credential service add-on - AND - IP-address-to- username mappings identified through User Mapping, GlobalProtect, or Authentication Policy and Captive Portal.	The firewall determines if the username and password a user submits matches the same user's corporate username and password. To do this, the firewall must able to match credential submissions to valid corporate usernames and passwords and verify that the username submitted maps to the IP address of the logged in user as follows: • **To detect corporate usernames and passwords**—The firewall retrieves a secure bit mask, called a *bloom filter*, from a Windows-based User-ID agent equipped with the User-ID credential service add-on. This add-on service scans your directory for usernames and password hashes and deconstructs them into a secure bit mask—the bloom filter—and delivers it to the Windows agent. The firewall retrieves the bloom filter from the Windows agent at regular intervals and, whenever it detects a user submitting credentials to a restricted category, it reconstructs the bloom filter and looks for a matching username and password hash. The firewall can only connect to one Windows-based User-ID agent running the User-ID credential service add-on. • **To verify that the credentials belong to the logged-in user**—The firewall looks for a mapping between the IP address of the logged-in user and the detected username in its IP-address-to-username mapping table. To learn more how the domain credential method works, and the requirements for enabling this type of detection, see Configure Credential Detection with the Windows-based User-ID Agent.

© 2018 Palo Alto Networks, Inc.

Configure Credential Detection with the Windows-based User-ID Agent

Domain Credential Filter detection enables the firewall to detect passwords submitted to web pages. This credential detection method requires the Windows-based User-ID agent and the User-ID credential service, an add-on to the User-ID agent, to be installed on a *read-only domain controller (RODC)*.



The Domain Credential Filter detection method is supported with the Windows-based User-ID agent only. You cannot use the PAN-OS integrated User-ID agent to configure this method of credential detection.

An RODC is a Microsoft Windows server that maintains a read-only copy of an Active Directory database that a domain controller hosts. When the domain controller is located at a corporate headquarters, for example, RODCs can be deployed in remote network locations to provide local authentication services. Installing the User-ID agent on an RODC can be useful for a few reasons: access to the domain controller directory is not required to enable credential detection and you can support credential detection for a limited or targeted set of users. Because the directory the RODC hosts is read-only, the directory contents remain secure on the domain controller.



Because you must install the Windows-based User-ID agent on the RODC for credential detection, as a best practice deploy a separate agent for this purpose. Do not use the User-ID agent installed on the RODC to map IP addresses to users.

After you install the User-ID agent on an RODC, the User-ID credential service runs in the background and scans the directory for the usernames and password hashes of group members that are listed in the RODC password replication policy (PRP)—you can define who you want to be on this list. The User-ID credential service then takes the collected usernames and password hashes and deconstructs the data into a type of bit mask called a *bloom filter*. Bloom filters are compact data structures that provide a secure method to check if an element (a username or a password hash) is a member of a set of elements (the sets of credentials you have approved for replication to the RODC). The User-ID credential service forwards the bloom filter to the Windows-based User-ID agent; the firewall retrieves the latest bloom filter from the User-ID agent at regular intervals and uses it to detect usernames and password hash submissions. Depending on your settings, the firewall then blocks, alerts, or allows on valid password submissions to web pages, or displays a response page to users warning them of the dangers of phishing, but allowing them to continue with the submission.

Throughout this process, the User-ID agent does not store or expose any password hashes, nor does it forward password hashes to the firewall. Once the password hashes are deconstructed into a bloom filter, there is no way to recover them.

STEP 1 | Configure User Mapping Using the Windows User-ID Agent.



To enable credential detection, you must install the Windows-based User-ID agent on an RODC (requires Microsoft Windows 2008r2 64 or later). Install a separate User-ID agent for this purpose.

Important items to remember when setting up User-ID to enable Domain Credential Filter detection:

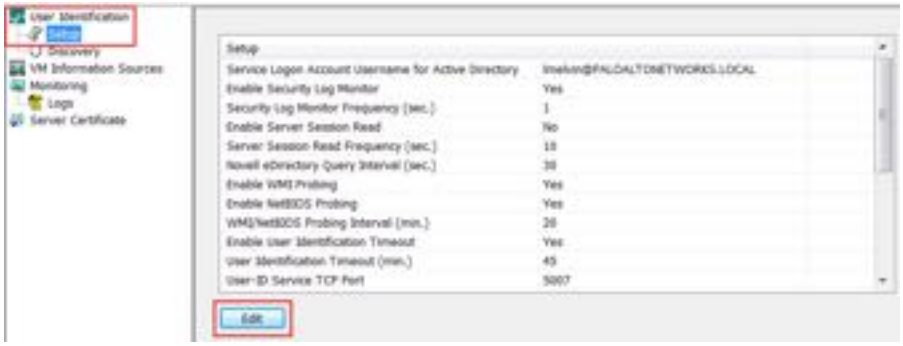
- Because the effectiveness of credential phishing detection is dependent on your RODC setup, make sure that you also review best practices and recommendations for RODC Administration.
- Download User-ID software updates:
  - User-ID Agent Windows installer—UaInstall-x.x.x-x.msi.

© 2018 Palo Alto Networks, Inc.

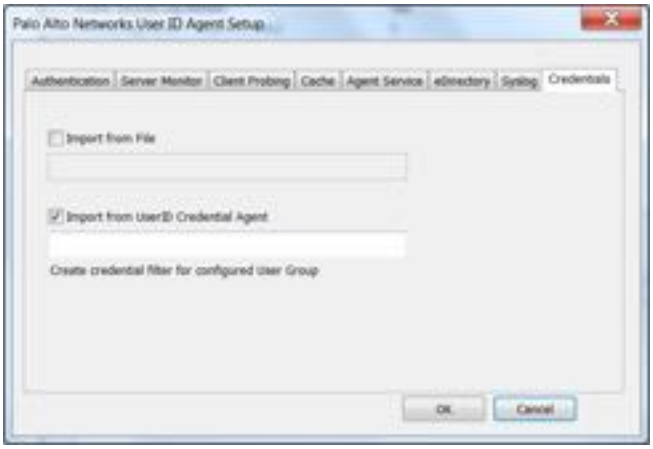
- User-ID Agent Credential Service Windows installer—UaCredInstall64-x.x.x-x.msi.
- Install the User-ID agent and the User Agent Credential service on an RODC using an account that has privileges to read Active Directory via LDAP (the User-ID agent also requires this privilege).

STEP 2 | Enable the User-ID agent and the User Agent Credential service (which runs in the background to scan permitted credentials) to share information.

1. On the RODC server, launch the User-ID Agent.
2. Select **Setup** and edit the Setup section.



3. Select the **Credentials** tab. This tab only displays if you have already installed the User-ID Agent Credential Service.



4. Select **Import from User-ID Credential Agent**. This enables the User-ID agent to import the bloom filter that the User-ID credential agent creates to represent users and the corresponding password hashes.
5. Click **OK**, **Save** your settings, and **Commit**.

STEP 3 | In the RODC directory, define the group of users for which you want to support credential submission detection.

- Confirm that the groups that should receive credential submission enforcement are added to the Allowed RODC Password Replication Group.
- Check that none of the groups in the Allowed RODC Password Replication Group are also in the Denied RODC Password Replication Group by default. Groups listed in both will not be subject to credential phishing enforcement.

STEP 4 | Continue to the next task.

Set Up Credential Phishing Prevention on the firewall.

© 2018 Palo Alto Networks, Inc.

Set Up Credential Phishing Prevention

After you have decided which of the Methods to Check for Corporate Credential Submissions you want to use, take the following steps to enable the firewall to detect when users submit corporate credentials to web pages and either alert on this action, block the credential submission, or require users to acknowledge the dangers of phishing before continuing with credential submission.

STEP 1 | If you have not done so already, Enable User-ID.

Each of the Methods to Check for Corporate Credential Submissions requires a different User-ID configuration to check for corporate credential submissions:

- If you plan to use the group mapping method, which detects whether a user is submitting a valid corporate username, Map Users to Groups.
- If you plan to use the IP user mapping method, which detects whether a user is submitting a valid corporate username and that username belongs to the logged-in user, Map IP Addresses to Users.
- If you plan to use the domain credential filter method, which detects whether a user is submitting a valid username and password and that those credentials belong to the logged in user, Configure Credential Detection with the Windows-based User-ID Agent and Map IP Addresses to Users.

STEP 2 | If you have not done so already, configure a best practice URL Filtering profile to ensure protection against URLs that have been observed hosting malware or exploitive content.

1. Select **Objects** > **Security Profiles** > **URL Filtering** and **Add** or modify a URL Filtering profile.
2. Block access to all known dangerous URL categories: malware, phishing, dynamic DNS, unknown, questionable, extremism, copyright-infringement, proxy-avoidance-and-anonymizers, and parked.

STEP 3 | Configure the URL Filtering profile to detect corporate credential submissions to websites that are in allowed URL categories.



The firewall automatically skips checking credential submissions for App-IDs associated with sites that have never been observed hosting malware or phishing content to ensure the best performance even if you enable checks in the corresponding category. The list of sites on which the firewall will skip credential checking is automatically updated via Application and Threat content updates.

1. Select **User Credential Detection**.
2. Select one of the Methods to Check for Corporate Credential Submissions to web pages from the **User Credential Detection** drop-down:
   - **Use IP User Mapping**—Checks for valid corporate username submissions and verifies that the login username maps to the source IP address of the session. To do this, the firewall matches the submitted username and source IP address of the session against its IP-address-to-username mapping table. To use this method you can use any of the user mapping methods described in Map IP Addresses to Users.
   - **Use Domain Credential Filter**—Checks for valid corporate usernames and password submissions verifies that the username maps to the IP address of the logged in user. See Configure Credential Detection with the Windows-based User-ID Agent for instructions on how to set up User-ID to enable this method.
   - **Use Group Mapping**—Checks for valid username submissions based on the user-to-group mapping table populated when you configure the firewall to Map Users to Groups.

     With group mapping, you can apply credential detection to **any** part of the directory, or specific group, such as groups like IT that have access to your most sensitive applications.

© 2018 Palo Alto Networks, Inc.

This method is prone to false positives in environments that do not have uniquely structured usernames. Because of this, you should only use this method to protect your high-value user accounts.

3. Set the **Valid Username Detected Log Severity** the firewall uses to log detection of corporate credential submissions. By default, the firewall logs these events as medium severity.

STEP 4 | Block (or alert) on credential submissions to allowed sites.

1. Select **Categories**.
2. For each Category to which **Site Access** is allowed, select how you want to treat **User Credential Submissions**:
   - **alert**—Allow users to submit credentials to the website, but generate a URL Filtering log each time a user submits credentials to 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, preventing the credential submission.
   - **continue**—Present the Anti Phishing Continue Page response page to users when they attempt to submit credentials. Users must select Continue on the response page to continue with the submission.
3. Select **OK** to save the URL Filtering profile.

STEP 5 | Apply the URL Filtering profile with the credential detection settings to your Security policy rules.

1. Select **Policies** > **Security** and **Add** or modify a Security policy rule.
2. On the **Actions** tab, set the **Profile Type** to **Profiles**.
3. Select the new or updated **URL Filtering** profile to attach it to the Security policy rule.
4. Select **OK** to save the Security policy rule.

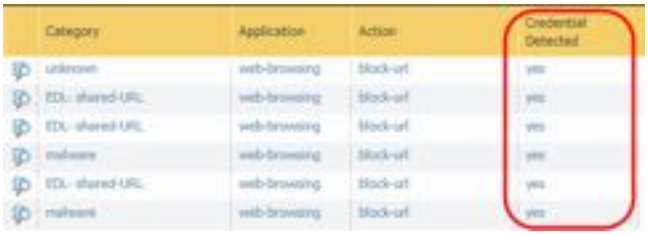
STEP 6 | **Commit** the configuration.

STEP 7 | Monitor credential submissions the firewall detects.

Select ACC > Hosts Visiting Malicious URLs to see the number of users who have visited malware and phishing sites.

Select **Monitor** > **Logs** > **URL Filtering**.

The new **Credential Detected** column indicates events where the firewall detected a HTTP post request that included a valid credential:



(To display this column, hover over any column header and click the arrow to select the columns you'd like to display).

Log entry details also indicate credential submissions:

© 2018 Palo Alto Networks, Inc.



STEP 8 | Validate and troubleshoot credential submission detection.

- Use the following CLI command to view credential detection statistics:

```
> show user credential-filter statistics
```

The output for this command varies depending on the method configured for the firewall to detect credential submissions. For example, if the Domain Credential Filter method is configured in any URL Filtering profile, a list of User-ID agents that have forwarded a bloom filter to the firewall is displayed, along with the number of credentials contained in the bloom filter.

- (Group Mapping method only) Use the following CLI command to view group mapping information, including the number of URL Filtering profiles with Group Mapping credential detection enabled and the usernames of group members that have attempted to submit credentials to a restricted site.

```
> show user group-mapping statistics
```

- (Domain Credential Filter method only) Use the following CLI command to see all Windows-based User-ID agents that are sending mappings to the firewall:

```
> show user user-id-agent state
```

The command output now displays bloom filter counts that include the number of bloom filter updates the firewall has received from each agent, if any bloom filter updates failed to process, and how many seconds have passed since the last bloom filter update.

- (Domain Credential Filter method only) The Windows-based User-ID agent displays log messages that reference BF (bloom filter) pushes to the firewall. In the User-ID agent interface, select **Monitoring** > **Logs**.

© 2018 Palo Alto Networks, Inc.

Share Threat Intelligence with Palo Alto Networks

Telemetry is the process of collecting and transmitting data for analysis. When you enable telemetry on the firewall, the firewall periodically collects and sends information that includes applications, threats, and device health to Palo Alto Networks. Sharing threat intelligence provides the following benefits:

- Enhanced vulnerability and spyware signatures delivered to you and other customers worldwide. For example, when a threat event triggers vulnerability or spyware signatures, the firewall shares the URLs associated with the threat with the Palo Alto Networks threat research team, so they can properly classify the URLs as malicious.
- Rapid testing and evaluation of experimental threat signatures with no impact to your network, so that critical threat prevention signatures can be released to all Palo Alto Networks customers faster.
- Improved accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control (C2) signatures, and WildFire.

Palo Alto Networks uses the threat intelligence extracted from telemetry to deliver these benefits to you and other Palo Alto Networks users. All Palo Alto Networks users benefit from the telemetry data shared by each user, making telemetry a community-driven approach to threat prevention. Palo Alto Networks does not share your telemetry data with other customers or third-party organizations.

- What Telemetry Data Does the Firewall Collect?
- Passive DNS Monitoring
- Enable Telemetry

What Telemetry Data Does the Firewall Collect?

The firewall collects and forwards different sets of telemetry data to Palo Alto Networks based on the Telemetry settings you enable. The firewall collects the data from fields in your log entries (see Log Types and Severity Levels); the log type and combination of fields vary based on the setting. Review the following table before you Enable Telemetry.

Setting	Description
Application Reports	The number and size of known applications by destination port, unknown applications by destination port, and unknown applications by destination IP address. The firewall generates these reports from Traffic logs and forwards them every 4 hours.
Threat Prevention Reports	Attacker information, the number of threats for each source country and destination port, and the correlation objects that threat events triggered.The firewall generates these reports from Threat logs and forwards them every 4 hours.
URL Reports	URLs with the following PAN-DB URL categories: malware, phishing, dynamic DNS, proxy-avoidance, questionable, parked, and unknown (URLs that PAN-DB has not yet categorized). The firewall generates these reports from URL Filtering logs. URL Reports also include PAN-DB statistics such as the version of the URL filtering database on the firewall and on the PAN-DB cloud, the number of URLs in those databases, and the number of URLs that the firewall

© 2018 Palo Alto Networks, Inc.

Setting	Description
	categorized. These statistics are based on the time that the firewall forwarded the URL Reports. The firewall forwards URL Reports every 4 hours.
File Type Identification Reports	Information about files that the firewall has blocked or allowed based on data filtering and file blocking settings. The firewall generates these reports from Data Filtering logs and forwards them every 4 hours.
Threat Prevention Data	Log data from threat events that triggered signatures that Palo Alto Networks is evaluating for efficacy. Threat Prevention Data provides Palo Alto Networks more visibility into your network traffic than other telemetry settings. When enabled, the firewall may collect information such as source or victim IP addresses. Enabling Threat Prevention Data also allows unreleased signatures that Palo Alto Networks is currently testing to run in the background. These signatures do not affect your security policy rules and firewall logs, and have no impact to your firewall performance. The firewall forwards Threat Prevention Data every 5 minutes.
Threat Prevention Packet Captures	Packet captures (if you have enabled your firewall to Take a Threat Packet Capture) of threat events that triggered signatures that Palo Alto Networks is evaluating for efficacy. Threat Prevention Packet Captures provide Palo Alto Networks more visibility into your network traffic than other telemetry settings. When enabled, the firewall may collect information such as source or victim IP addresses. The firewall forwards Threat Prevention Packet Captures every 5 minutes.
Product Usage Statistics	Back traces of firewall processes that have failed, as well as information about the firewall status. Back traces outline the execution history of the failed processes. These reports include details about the firewall model and the PAN-OS and content release versions installed on your firewall. The firewall forwards Product Usage Statistics every 5 minutes.
Passive DNS Monitoring	Domain-to-IP address mappings based on firewall traffic. When you enable Passive DNS Monitoring, the firewall acts as a passive DNS sensor and send DNS information to Palo Alto Networks for analysis. The firewall forwards data from Passive DNS Monitoring in 1 MB batches.

Passive DNS Monitoring

Passive DNS monitoring enables the firewall to act as a passive DNS sensor and send DNS information to Palo Alto Networks for analysis to improve threat intelligence and threat prevention capabilities. The data collected includes non-recursive DNS query (that is, the web browser sends a query to a DNS server to translate a domain to an IP address, and the server returns a response without querying other DNS servers) and response packet payloads. See DNS Overview for more background information about DNS.

The threat intelligence that the firewall collects from passive DNS monitoring consists solely of domain-to-IP address mappings. Palo Alto Networks retains no record of the source of this data and does not have the ability to associate it with the submitter at a future date. The Palo Alto Networks threat research team uses

© 2018 Palo Alto Networks, Inc.

passive DNS information to gain insight into malware propagation and evasion techniques that abuse the DNS system. Information gathered through this data collection is used to improve PAN-DB URL category and DNS-based C2 signature accuracy and WildFire malware detection.

The firewall forwards DNS responses only when the following requirements are met:

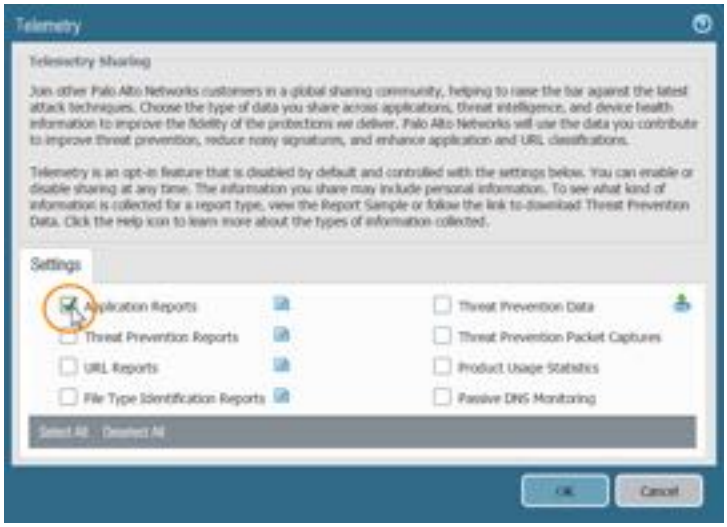
- DNS response bit is set
- DNS truncated bit is not set
- DNS recursive bit is not set
- DNS response code is 0 or 3 (NX)
- DNS question count bigger than 0
- DNS Answer RR count is bigger than 0 or if it is 0, the flags need to be 3 (NX)
- DNS query record type are A, NS, CNAME, AAAA, MX

Enable Telemetry

When you enable telemetry, you define what data the firewall collects and shares with Palo Alto Networks. For some telemetry settings, you can preview what the data that your firewall sends will look like before committing. The firewall uses the **Palo Alto Networks Services** service route to send the data you share from telemetry to Palo Alto Networks.

STEP 1 | Select **Device** > **Setup** > **Telemetry**, and edit the Telemetry settings.

STEP 2 | Select the telemetry data you want to share with Palo Alto Networks. For more specific descriptions of this data, see What Telemetry Data Does the Firewall Collect? By default, all telemetry settings are disabled.



To enable Threat Prevention Packet Captures, you must also enable Threat Prevention Data.

STEP 3 | Open a report sample ( ) to view the type of data that the firewall collects for Application Reports, Threat Prevention Reports, URL Reports, and File Type Identification Reports.

The report sample, formatted in XML, is based on your firewall activity in the first 4 hours since you first viewed the report sample. A report sample does not display any entries if the firewall did not find any matching traffic for the report. The firewall only collects new information for a report sample when you restart the firewall and open a report sample.

© 2018 Palo Alto Networks, Inc.

The figure below shows a report sample for Threat Prevention Reports:



Application Reports, Threat Prevention Reports, URL Reports, and File Type Identification Reports each consist of multiple reports. In the report sample, **Type** describes the name of a report. **Aggregate** lists the log fields that the firewall collects for the report (refer to Syslog Field Descriptions to determine the name of the fields as they appear in the firewall logs). **Values** indicates the units of measure used in the report (for example, the value **count** for the **Attackers (threat)** report refers to the number of times the firewall detected a threat associated with a particular threat ID).

STEP 4 | View the type of data that the firewall collects for Product Usage Statistics.

Enter the following operational CLI command: **`show system info`**

STEP 5 | Click **OK** and **Commit** your changes.

STEP 6 | If you enabled Threat Prevention Data and Threat Prevention Packet Captures, view the data that the firewall collected.

1. Edit the Telemetry settings.
2. Click **Download Threat Prevention Data** ( ) to download a tarball file (.tar.gz) with the most recent 100 folders of data that the firewall collected for Threat Prevention Data and Threat Prevention Packet Captures. If you never enabled these settings or if you enabled them but no threat events have matched the conditions for these settings, the firewall does not generate a file and instead returns an error message.

There is currently no way to view the DNS information that the firewall collects through passive DNS monitoring.

© 2018 Palo Alto Networks, Inc.

Use DNS Queries to Identify Infected Hosts on the Network

The DNS sinkhole action in Anti-Spyware profiles enables the firewall to forge a response to a DNS query for a known malicious domain or to a custom domain so that you can identify hosts on your network that have been infected with malware. By default, DNS queries to any domain included in the Palo Alto Networks DNS signatures list is sinkholed to a Palo Alto Networks server IP address. The following topics provide details on how to enable DNS sinkholing for custom domains and how to identify infected hosts.

- DNS Sinkholing
- Configure DNS Sinkholing for a List of Custom Domains
- Configure the Sinkhole IP Address to a Local Server on Your Network
- Identify Infected Hosts

DNS Sinkholing

DNS sinkholing helps you to identify infected hosts on the protected network using DNS traffic in situations where the firewall cannot see the infected client's DNS query (that is, the firewall cannot see the originator of the DNS query). In a typical deployment where the firewall is north of the local DNS server, the threat log will identify the local DNS resolver as the source of the traffic rather than the actual infected host. Sinkholing malware DNS queries solves this visibility problem by forging responses to the client host queries directed at malicious domains, so that clients attempting to connect to malicious domains (for command-and-control, for example) will instead attempt to connect to a default Palo Alto Networks sinkhole IP address, or to a user-defined IP address as illustrated in Configure DNS Sinkholing for a List of Custom Domains. Infected hosts can then be easily identified in the traffic logs because any host that attempts to connect to the sinkhole IP address is most likely infected with malware.

If you want to enable DNS sinkholing for Palo Alto Networks DNS signatures, attach the default Anti-Spyware profile to a security policy rule (see Set Up Antivirus, Anti-Spyware, and Vulnerability Protection). DNS queries to any domain included in the Palo Alto Networks DNS signatures will be resolved to the default Palo Alto Networks sinkhole IP address. The IP addresses currently are IPv4—71.19.152.112 and a loopback address IPv6 address—::1. These address are subject to change and can be updated with content updates.

© 2018 Palo Alto Networks, Inc.

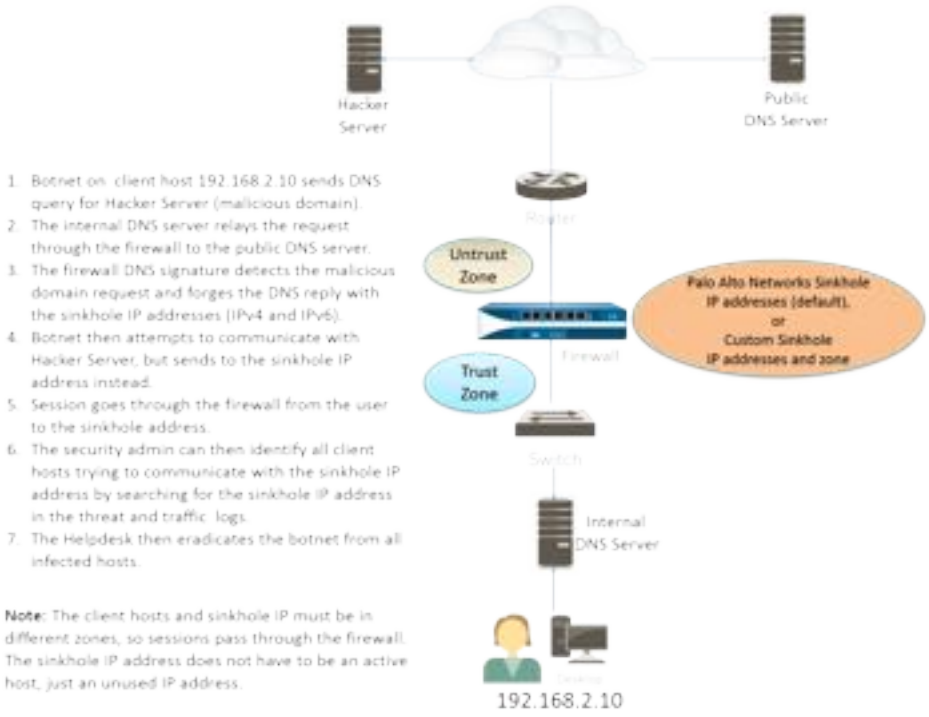


Figure 7: DNS Sinkholing Example

Configure DNS Sinkholing for a List of Custom Domains

To enable DNS Sinkholing for a custom list of domains, you must create an External Dynamic List that includes the domains, enable the sinkhole action in an Anti-Spyware profile and attach the profile to a security policy rule. When a client attempts to access a malicious domain in the list, the firewall forges the destination IP address in the packet to the default Palo Alto Networks server or to a user-defined IP address for sinkholing.

For each custom domain included in the external dynamic list, the firewall generates DNS-based spyware signatures. The signature is named Custom Malicious DNS Query <domain name>, and is of type spyware with medium severity; each signature is a 24-byte hash of the domain name.

Each firewall model supports a maximum of 50,000 domain names total in one or more external dynamic lists but no maximum limit is enforced for any one list.

STEP 1 | Enable DNS sinkholing for the custom list of domains in an external dynamic list.

1. Select **Objects** > **Security Profiles** > **Anti-Spyware**.
2. Modify an existing profile, or select one of the existing default profiles and clone it.
3. **Name** the profile and select the **DNS Signatures** tab.
4. Click **Add** and select **External Dynamic Lists** in the drop-down.

   If you have already created an external dynamic list of type: Domain List, you can select it from here. The drop-down does not display external dynamic lists of type URL or IP Address that you may have created.
5. Configure the external dynamic list from the Anti-Spyware profile (see Configure the Firewall to Access an External Dynamic List). The **Type** is preset to **Domain List**.

© 2018 Palo Alto Networks, Inc.

6. (Optional) In the **Packet Capture** drop-down, select **single-packet** to capture the first packet of the session or **extended-capture** to set between 1-50 packets. You can then use the packet captures for further analysis.

STEP 2 | Verify the sinkholing settings on the Anti-Spyware profile.

1. On the **DNS Signatures** tab, verify that the **Action on DNS Queries** is **sinkhole**.
2. In the Sinkhole section, verify that **Sinkhole** is enabled. For your convenience, the default Sinkhole IP address is set to access a Palo Alto Networks server. Palo Alto Networks can automatically refresh this IP address through content updates.

   If you want to modify the **Sinkhole IPv4** or **Sinkhole IPv6** address to a local server on your network or to a loopback address, see Configure the Sinkhole IP Address to a Local Server on Your Network.



3. Click **OK** to save the Anti-Spyware profile.

STEP 3 | Attach the Anti-Spyware profile to a Security policy rule.

1. Select **Policies** > **Security**.
2. On the **Actions** tab, select the **Log at Session Start** check box to enable logging.
3. In the Profile Setting section, click the **Profile Type** drop-down to view all **Profiles**. From the **Anti-Spyware** drop-down and select the new profile.
4. Click **OK** to save the policy rule.

STEP 4 | Test that the policy action is enforced.

1. View External Dynamic List Entries that belong to the domain list, and access a domain from the list.
2. To monitor the activity on the firewall:
   1. Select **ACC** and add a URL Domain as a global filter to view the Threat Activity and Blocked Activity for the domain you accessed.
   2. Select **Monitor** > **Logs** > **Threat** and filter by `(action eq sinkhole)` to view logs on sinkholed domains.

STEP 5 | Verify whether entries in the external dynamic list are ignored or skipped.

Use the following CLI command on the firewall to review the details about the list.

```
request system external-list show type domain name <list_name>
```

For example:

```
request system external-list show type domain name
```


© 2018 Palo Alto Networks, Inc.

```
My_List_of_Domains_2015
vsys1/EBLDomain:
Next update at : Thu May 21 10:15:39 2015
Source : https://1.2.3.4/My_List_of_Domains_2015
Referenced : Yes
Valid : Yes
Number of entries : 3
domains:www.example.com
baddomain.com
qqq.abcedfg.com
```

STEP 6 | (Optional) Retrieve the external dynamic list on-demand.

To force the firewall to retrieve the updated list on-demand instead of at the next refresh interval (the **Repeat** frequency you defined for the external dynamic list), use the following CLI command:

```
request system external-list refresh type domain name  <list_name>
```

As an alternative, you can use the firewall interface to Retrieve an External Dynamic List from the Web Server.

Configure the Sinkhole IP Address to a Local Server on Your Network

By default, sinkholing is enabled for all Palo Alto Networks DNS signatures, and the sinkhole IP address is set to access a Palo Alto Networks server. Use the instructions in this section if you want to set the sinkhole IP address to a local server on your network.

You must obtain both an IPv4 and IPv6 address to use as the sinkhole IP addresses because malicious software may perform DNS queries using one or both of these protocols. The DNS sinkhole address must be in a different zone than the client hosts to ensure that when an infected host attempts to start a session with the sinkhole IP address, it will be routed through the firewall.

The sinkhole addresses must be reserved for this purpose and do not need to be assigned to a physical host. You can optionally use a honey-pot server as a physical host to further analyze the malicious traffic.

The configuration steps that follow use the following example DNS sinkhole addresses:

IPv4 DNS sinkhole address—10.15.0.20

IPv6 DNS sinkhole address—fd97:3dec:4d27:e37c:5:5:5:5

STEP 1 | Configure the sinkhole interface and zone.

Traffic from the zone where the client hosts reside must route to the zone where the sinkhole IP address is defined, so traffic will be logged.

Use a dedicated zone for sinkhole traffic, because the infected host will be sending traffic to this zone.

1. Select **Network** > **Interfaces** and select an interface to configure as your sinkhole interface.
2. In the **Interface Type** drop-down, select **Layer3**.

© 2018 Palo Alto Networks, Inc.

3. To add an IPv4 address, select the **IPv4** tab and select **Static** and then click **Add**. In this example, add 10.15.0.20 as the IPv4 DNS sinkhole address.
4. Select the **IPv6** tab and click **Static** and then click **Add** and enter an IPv6 address and subnet mask. In this example, enter fd97:3dec:4d27:e37c::/64 as the IPv6 sinkhole address.
5. Click **OK** to save.
6. To add a zone for the sinkhole, select **Network** > **Zones** and click **Add**.
7. Enter zone **Name**.
8. In the **Type** drop-down select **Layer3**.
9. In the **Interfaces** section, click **Add** and add the interface you just configured.
10. Click **OK**.

STEP 2 | Enable DNS sinkholing.

By default, sinkholing is enabled for all Palo Alto Networks DNS signatures. To change the sinkhole address to your local server, see Step Verify the sinkholing settings on the Anti-Spyware profile. in Configure DNS Sinkholing for a List of Custom Domains.

STEP 3 | Edit the security policy rule that allows traffic from client hosts in the trust zone to the untrust zone to include the sinkhole zone as a destination and attach the Anti-Spyware profile.

Editing the Security policy rule(s) that allows traffic from client hosts in the trust zone to the untrust zone ensures that you are identifying traffic from infected hosts. By adding the sinkhole zone as a destination on the rule, you enable infected clients to send bogus DNS queries to the DNS sinkhole.

1. Select **Policies** > **Security**.
2. Select an existing rule that allows traffic from the client host zone to the untrust zone.
3. On the **Destination** tab, **Add** the Sinkhole zone. This allows client host traffic to flow to the sinkhole zone.
4. On the **Actions** tab, select the **Log at Session Start** check box to enable logging. This will ensure that traffic from client hosts in the Trust zone will be logged when accessing the Untrust or Sinkhole zones.
5. In the **Profile Setting** section, select the **Anti-Spyware** profile in which you enabled DNS sinkholing.
6. Click **OK** to save the Security policy rule and then **Commit**.

STEP 4 | To confirm that you will be able to identify infected hosts, verify that traffic going from the client host in the Trust zone to the new Sinkhole zone is being logged.

In this example, the infected client host is 192.168.2.10 and the Sinkhole IPv4 address is 10.15.0.20.

1. From a client host in the trust zone, open a command prompt and run the following command:

```
C:\>ping <sinkhole address>
```

The following example output shows the ping request to the DNS sinkhole address at 10.15.0.2 and the result, which is `Request timed out` because in this example the sinkhole IP address is not assigned to a physical host:

```
C:\>ping 10.15.0.20
Pinging 10.15.0.20 with 32 bytes of data:
Request timed out.
Request timed out.
Ping statistics for 10.15.0.20:
Packets: Sent = 4, Received = 0, Lost = 4 (100% loss)
```


© 2018 Palo Alto Networks, Inc.

2. On the firewall, select **Monitor** > **Logs** > **Traffic** and find the log entry with the Source 192.168.2.10 and Destination 10.15.0.20. This will confirm that the traffic to the sinkhole IP address is traversing the firewall zones.



You can search and/or filter the logs and only show logs with the destination 10.15.0.20. To do this, click the IP address (10.15.0.20) in the Destination column, which will add the filter (addr.dst in 10.15.0.20) to the search field. Click the Apply Filter icon to the right of the search field to apply the filter.

STEP 5 | Test that DNS sinkholing is configured properly.

You are simulating the action that an infected client host would perform when a malicious application attempts to call home.

1. Find a malicious domain that is included in the firewall's current Antivirus signature database to test sinkholing.
   1. Select **Device** > **Dynamic Updates** and in the **Antivirus** section click the **Release Notes** link for the currently installed antivirus database. You can also find the antivirus release notes that list the incremental signature updates under Dynamic Updates on the Palo Alto Networks support site.
   2. In the second column of the release note, locate a line item with a domain extension (for example, .com, .edu, or .net). The left column will display the domain name. For example, Antivirus release 1117-1560, includes an item in the left column named "tbsbana" and the right column lists "net".

      The following shows the content in the release note for this line item:

```
conficker:tbsbana 1
variants: net
```

2. From the client host, open a command prompt.
3. Perform an NSLOOKUP to a URL that you identified as a known malicious domain.

   For example, using the URL `track.bidtrk.com`:

```
C:\>nslookup
 track.bidtrk.com
 Server: my-local-dns.local
 Address: 10.0.0.222
 Non-authoritative answer:
 Name: track.bidtrk.com.org
 Addresses: fd97:3dec:4d27:e37c:5:5:5:510.15.0.20
```

   In the output, note that the NSLOOKUP to the malicious domain has been forged using the sinkhole IP addresses that we configured (10.15.0.20). Because the domain matched a malicious DNS signature, the sinkhole action was performed.
4. Select **Monitor** > **Logs** > **Threat** and locate the corresponding threat log entry to verify that the correct action was taken on the NSLOOKUP request.
5. Perform a ping to **`track.bidtrk.com`**, which will generate network traffic to the sinkhole address.

Identify Infected Hosts

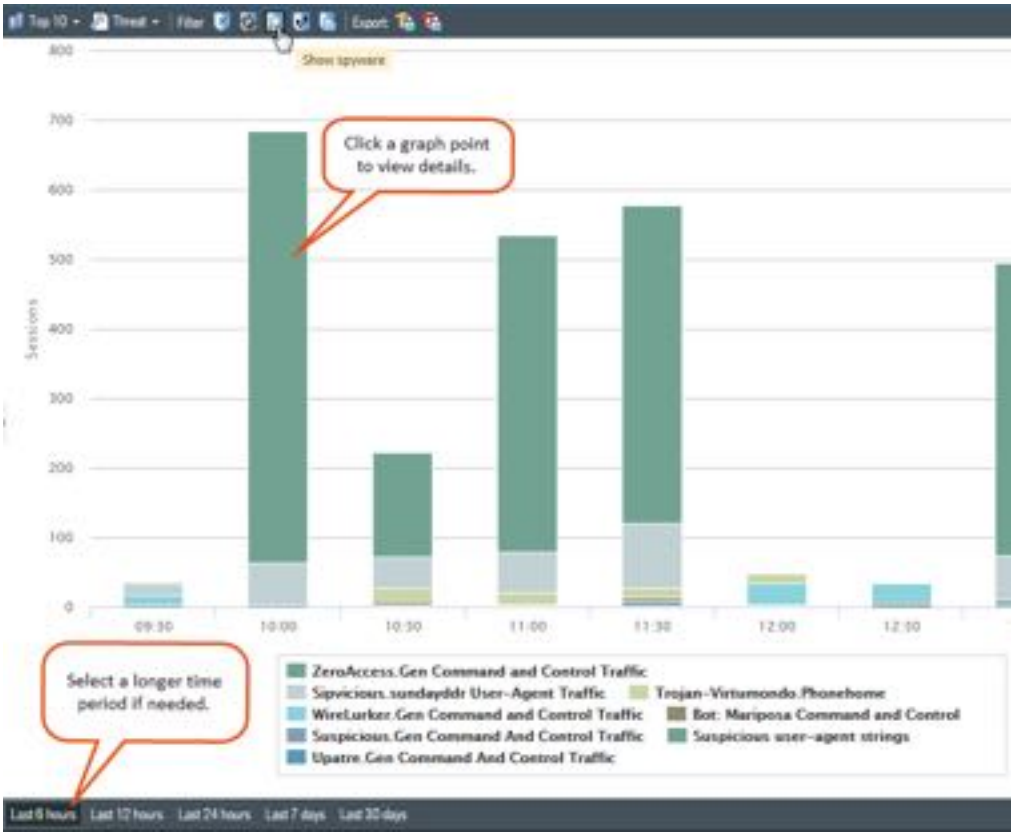
After you have configured DNS sinkholing and verified that traffic to a malicious domain goes to the sinkhole address, you should regularly monitor traffic to the sinkhole address, so that you can track down the infected hosts and eliminate the threat.

- Use App Scope to identify infected client hosts.

© 2018 Palo Alto Networks, Inc.

1. Select **Monitor** > **App Scope** and select **Threat Monitor**.
2. Click the **Show spyware** button along the top of the display page.
3. Select a time range.

   The following screenshot shows three instances of Suspicious DNS queries, which were generated when the test client host performed an NSLOOKUP on a known malicious domain. Click the graph to see more details about the event.



- Configure a custom report to identify all client hosts that have sent traffic to the sinkhole IP address, which is 10.15.0.20 in this example.

  Forward to an SNMP manager, Syslog server and/or Panorama to enable alerts on these events.

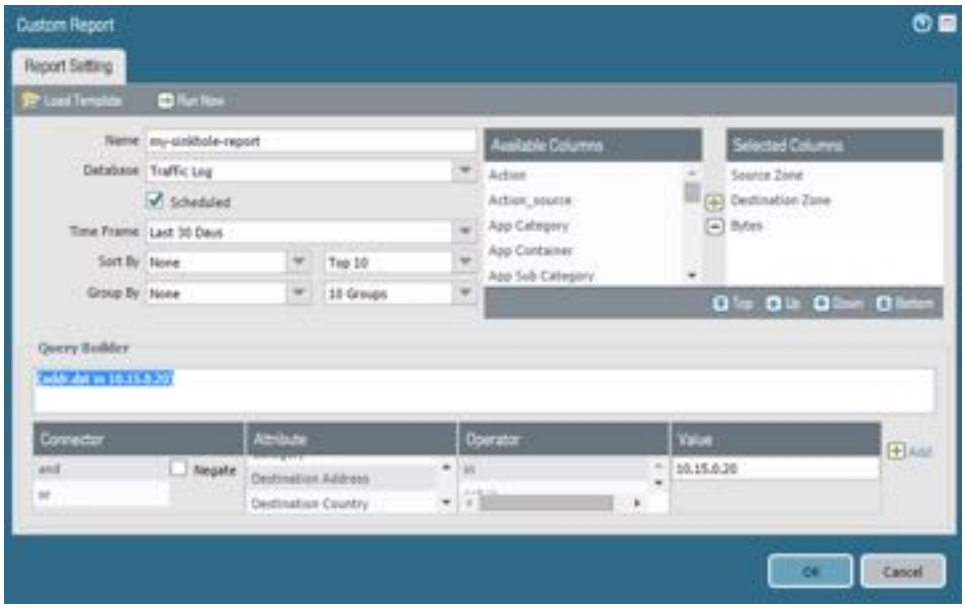
  In this example, the infected client host performed an NSLOOKUP to a known malicious domain that is listed in the Palo Alto Networks DNS Signature database. When this occurred, the query was sent to the local DNS server, which then forwarded the request through the firewall to an external DNS server. The firewall security policy with the Anti-Spyware profile configured matched the query to the DNS Signature database, which then forged the reply using the sinkhole address of 10.15.0.20 and fd97:3dec:4d27:e37c:5:5:5:5. The client attempts to start a session and the traffic log records the activity with the source host and the destination address, which is now directed to the forged sinkhole address.

  Viewing the traffic log on the firewall allows you to identify any client host that is sending traffic to the sinkhole address. In this example, the logs show that the source address 192.168.2.10 sent the malicious DNS query. The host can then be found and cleaned. Without the DNS sinkhole option, the administrator would only see the local DNS server as the system that performed the query and would

© 2018 Palo Alto Networks, Inc.

not see the client host that is infected. If you attempted to run a report on the threat log using the action "Sinkhole", the log would show the local DNS server, not the infected host.

1. Select **Monitor** > **Manage Custom Reports**.
2. Click **Add** and **Name** the report.
3. Define a custom report that captures traffic to the sinkhole address as follows:
   - **Database**—Select **Traffic Log**.
   - **Scheduled**—Enable **Scheduled** and the report will run every night.
   - **Time Frame**—30 days
   - **Selected Columns**—Select **Source address** or **Source User** (if you have User-ID configured), which will identify the infected client host in the report, and **Destination address**, which will be the sinkhole address.
   - In the section at the bottom of the screen, create a custom query for traffic to the sinkhole address (10.15.0.20 in this example). You can either enter the destination address in the **Query Builder** window (**addr.dst in 10.15.0.20**) or select the following in each column and click **Add**: Connector = and, Attribute = Destination Address, Operator = in, and Value = 10.15.0.20. Click **Add** to add the query.



4. Click **Run Now** to run the report. The report will show all client hosts that have sent traffic to the sinkhole address, which indicates that they are most likely infected. You can now track down the hosts and check them for spyware.



5. To view scheduled reports that have run, select **Monitor** > **Reports**.

© 2018 Palo Alto Networks, Inc.

Monitor Blocked IP Addresses

The firewall maintains a block list of source IP addresses that it's blocking. When the firewall blocks a source IP address, such as when you configure either of the following policy rules, the firewall blocks that traffic in hardware before those packets use CPU or packet buffer resources:

- A classified DoS Protection policy rule with the action to **Protect** (a classified DoS Protection policy specifies that incoming connections match a source IP address, destination IP address, or source and destination IP address pair, and is associated with a Classified DoS Protection profile, as described in DoS Protection Against Flooding of New Sessions)
- A Security Policy rule that uses a Vulnerability Protection profile

Hardware IP address blocking is supported on PA-3060 firewalls, PA-3050 firewalls, and PA-5000 Series, PA-5200 Series, and PA-7000 Series firewalls.

You can view the block list, get detailed information about an IP address on the block list, or view counts of addresses that hardware and software are blocking. You can delete an IP address from the list if you think it shouldn't be blocked. You can change the source of detailed information about addresses on the list. You can also change how long hardware blocks IP addresses.

- View block list entries.
  1. Select **Monitor** > **Block IP List**.

     Entries on the block list indicate in the Type column whether they were blocked by hardware (hw) or software (sw).
  2. View at the bottom of the screen:
     - Count of **Total Blocked IPs** out of the number of blocked IP addresses the firewall supports.
     - Percentage of the block list the firewall has used.
  3. To filter the entries displayed, select a value in a column (which creates a filter in the **Filters** field) and Apply Filter ( ). Otherwise, the firewall displays the first 1,000 entries.
  4. Enter a **Page** number or click the arrows at the bottom of the screen to advance through pages of entries.
  5. To view details about an address on the block list, hover over a Source IP address and click the down arrow link. Click the **Who Is** link, which displays Network Solutions Who Is information about the address.



- Delete block list entries.

  You might want to delete an entry if you determine an IP address shouldn't be blocked. You should then revise the policy rule that caused the firewall to block the address.

  1. Select **Monitor** > **Block IP List**.

© 2018 Palo Alto Networks, Inc.

2. Select one or more entries and click **Delete**.
3. (Optional) Select **Clear All** to remove all entries from the list.

- Disable or re-enable hardware IP address blocking for troubleshooting purposes.



While hardware IP address blocking is disabled, the firewall still performs any software IP address blocking you have configured.

```
> set system setting hardware-acl-blocking [enable | disable]
```



Leave hardware IP address blocking enabled unless Palo Alto Networks technical support asks you to disable it, for example, if they are debugging a traffic flow.

- Tune the number of seconds that IP addresses blocked by hardware remain on the block list (range is 1-3,600; default is 1).

```
> set system setting hardware-acl-blocking duration <seconds>
```



Maintain a shorter duration for hardware block list entries than software block list entries to reduce the likelihood of exceeding the blocking capacity of the hardware.

- Change the default website for finding more information about an IP address from Network Solutions Who Is to a different website.

```
# set deviceconfig system ip-address-lookup-url <url>
```

- View counts of source IP addresses blocked by hardware and software, for example to see the rate of an attack.

  View the total sum of IP address entries on the hardware block table and block list (blocked by hardware and software):

```
> show counter global name flow_dos_blk_num_entries
```

  View the count of IP address entries on the hardware block table that were blocked by hardware:

```
> show counter global name flow_dos_blk_hw_entries
```

  View the count of IP address entries on the block list that were blocked by software:

```
> show counter global name flow_dos_blk_sw_entries
```

- View block list information per slot on a PA-7000 Series firewall.

```
> show dos-block-table software filter slot <slot-number>
```


© 2018 Palo Alto Networks, Inc.

Learn More About and Assess Threats

Features of Threat Vault and AutoFocus are integrated into the firewall to provide visibility into the nature of the threats the firewall detects and to give a more complete picture of how an artifact fits into your organization's network traffic (an artifact is property, activity, or behavior associated with a file, email link, or session). These features dually allow you get immediate, contextual information about a threat or to seamlessly shift your threat investigation from the firewall to the Threat Vault and AutoFocus.



Additionally, you can use threat categories—which classify types of threat events—to narrow your view into a certain type of threat activity or to build custom reports.

- Assess Firewall Artifacts with AutoFocus
- Learn More About Threat Signatures
- Monitor Activity and Create Custom Reports Based on Threat Categories

Assess Firewall Artifacts with AutoFocus

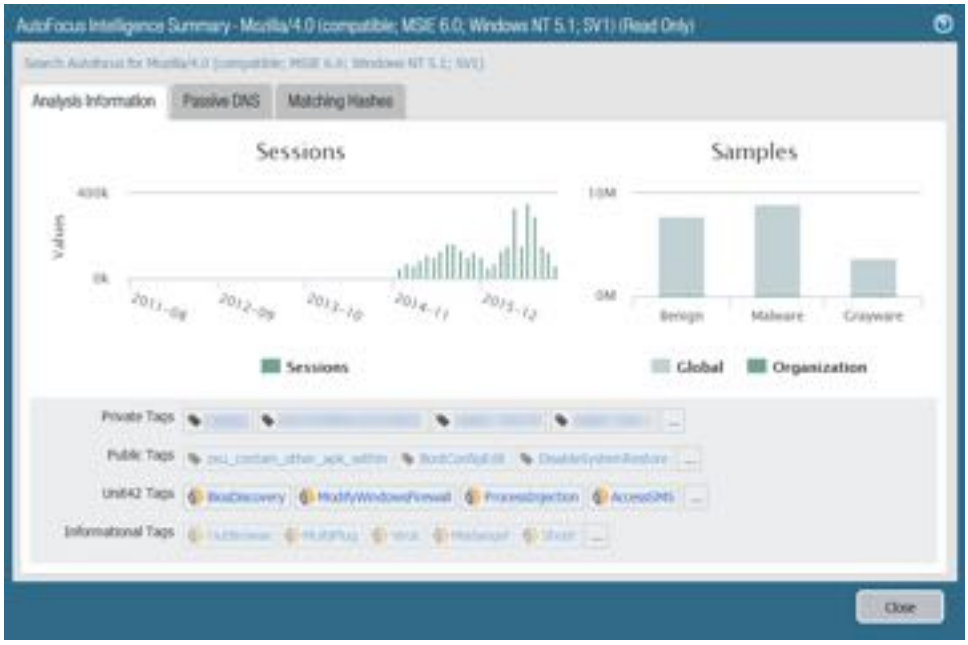
Use the AutoFocus Intelligence Summary for an artifact to assess its pervasiveness in your network and the threats associated with it.

- AutoFocus Intelligence Summary
- View and Act on AutoFocus Intelligence Summary Data

AutoFocus Intelligence Summary

The AutoFocus Intelligence Summary offers a centralized view of information about an artifact that AutoFocus has extracted from threat intelligence gathered from other AutoFocus users, WildFire, the PAN-DB URL filtering database, Unit 42, and open-source intelligence.

© 2018 Palo Alto Networks, Inc.



AutoFocus Intelligence Summary	
Analysis Information	The Analysis Information tab displays the following information: • Sessions—The number of sessions logged in your firewall(s) in which the firewall detected samples associated with the artifact. • Samples—A comparison of organization and global samples associated with the artifact and grouped by WildFire verdict (benign, malware, or grayware). *Global* refers to samples from all WildFire submissions, while *organization* refers only to samples submitted to WildFire by your organization. • Matching Tags—The AutoFocus tags matched to the artifact. AutoFocus Tags indicate whether an artifact is linked to malware or targeted attacks.
Passive DNS	The Passive DNS tab displays passive DNS history that includes the artifact. This passive DNS history is based on global DNS intelligence in AutoFocus; it is not limited to the DNS activity in your network. Passive DNS history consists of: • The domain request • The DNS request type • The IP address or domain to which the DNS request resolved (private IP addresses are not displayed) • The number of times the request was made • The date and time the request was first seen and last seen
Matching Hashes	The Matching Hashes tab displays the 5 most recently detected matching samples. Sample information includes: • The SHA256 hash of the sample • The sample file type • The date and time that WildFire analyzed a sample and assigned a WildFire verdict to it

© 2018 Palo Alto Networks, Inc.

AutoFocus Intelligence Summary	
	• The WildFire verdict for the sample • The date and time that WildFire updated the WildFire verdict for the sample (if applicable)

View and Act on AutoFocus Intelligence Summary Data

Interact with the AutoFocus Intelligence Summary to display more information about an artifact or extend your artifact research to AutoFocus. AutoFocus tags reveal if the artifact is associated with certain types of malware or malicious behavior.

STEP 1 | Confirm that the firewall is connected to AutoFocus.

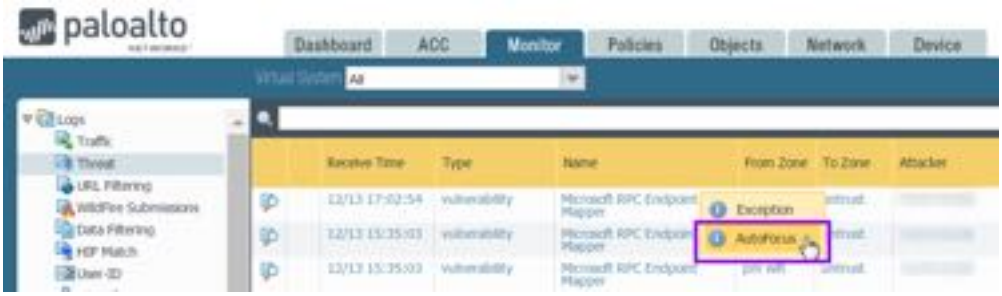
Enable AutoFocus Threat Intelligence on the firewall (active AutoFocus subscription required).

STEP 2 | Find artifacts to investigate.

You can view an AutoFocus Intelligence Summary for artifacts when you:

- View Logs (Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering, and Unified logs only).
- View External Dynamic List Entries.

STEP 3 | Hover over an artifact to open the drop-down, and click **AutoFocus**.



The AutoFocus Intelligence Summary is only available for the following types of artifacts:

IP address

URL

Domain

User agent

Threat name (only for threats of the subtypes virus and wildfire-virus)

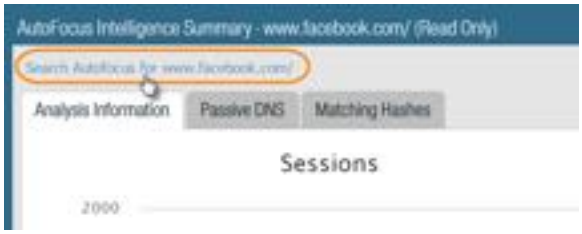
Filename

SHA-256 hash

STEP 4 | Launch an AutoFocus search for the artifact for which you opened the AutoFocus Intelligence Summary.

Click the **Search AutoFocus for...** link at the top of the AutoFocus Intelligence Summary window. The search results include all samples associated with the artifact. Toggle between the **My Samples** and **All Samples** tabs and compare the number of samples to determine the pervasiveness of the artifact in your organization.

© 2018 Palo Alto Networks, Inc.



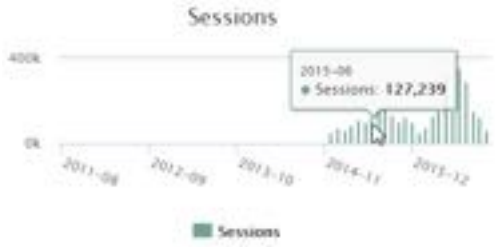
STEP 5 | Launch an AutoFocus search for other artifacts in the AutoFocus Intelligence Summary.

Click on the following artifacts to determine their pervasiveness in your organization:

- WildFire verdicts in the Analysis Information tab
- URLs and IP addresses in the Passive DNS tab
- The SHA256 hashes in the Matching Hashes tab

STEP 6 | View the number of sessions associated with the artifact in your organization per month.

Hover over the session bars.



STEP 7 | View the number of samples associated with the artifact by scope and WildFire verdict.

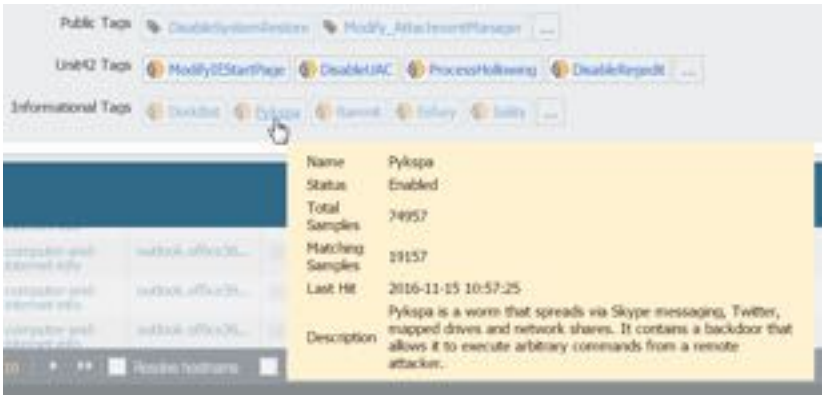
Hover over the samples bars.



STEP 8 | View more details about matching AutoFocus. tags.

Hover over a matching tag to view the tag description and other tag details.

© 2018 Palo Alto Networks, Inc.



STEP 9 | View other samples associated with a matching tag.

Click a matching tag to launch an AutoFocus search for that tag. The search results include all samples matched to the tag.

Unit 42 tags identify threats and campaigns that pose a direct security risk. Click on a Unit 42 matching tag to see how many samples in your network are associated with the threat the tag identifies.

STEP 10 | Find more matching tags for an artifact.

Click the ellipsis (...) to launch an AutoFocus search for the artifact. The Tags column in the search results displays more matching tags for the artifact, which give you an idea of other malware, malicious behavior, threat actors, exploits, or campaigns where the artifact is commonly detected.



Learn More About Threat Signatures

Firewall Threat logs record all threats the firewall detects based on threat signatures (Set Up Antivirus, Anti-Spyware, and Vulnerability Protection) and the ACC displays an overview of the top threats on your network. Each event the firewall records includes an ID that identifies the associated threat signature.

You can use the threat ID found with a Threat log or ACC entry to:

- Easily check if a threat signature is configured as an exception to your security policy (Create Threat Exceptions).
- Find the latest Threat Vault information about a specific threat. Because the Threat Vault is integrated with the firewall, you can view threat details directly in the firewall context or launch a Threat Vault search in a new browser window for a threat the firewall logged.

STEP 1 | Confirm the firewall is connected to the Threat Vault.

Select **Device** > **Setup** > **Management** and edit the **Logging and Reporting** setting to **Enable Threat Vault Access**. Threat vault access is enabled by default.

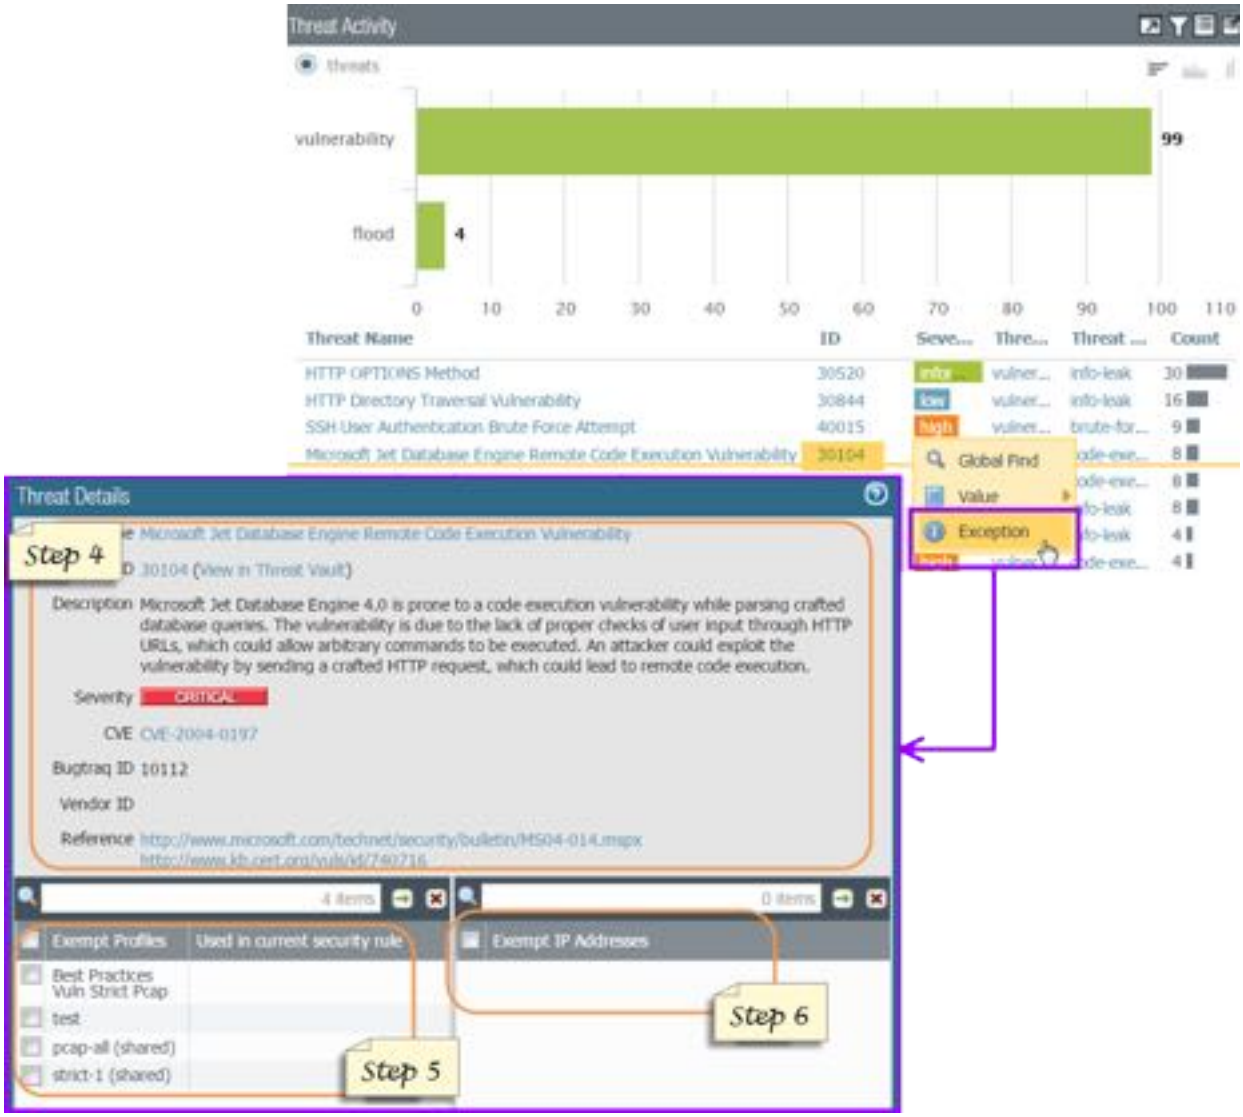
STEP 2 | Find the threat ID for threats the firewall detects.

© 2018 Palo Alto Networks, Inc.

- To see each threat event the firewall detects based on threat signatures, select **Monitor** > **Logs** > **Threat**. You can find the ID for a threat entry listed in the ID column, or select the log entry to view log details, including the Threat ID.
- To see an overview of top threats on the network, select **ACC** > **Threat Activity** and take a look at the Threat Activity widget. The ID column displays the threat ID for each threat displayed.
- To see details for threats that you can configure as threat exceptions (meaning, the firewall enforces the threat differently than the default action defined for the threat signature), select **Objects** > **Security Profiles** > **Anti-Spyware/Vulnerability Protection**. **Add** or modify a profile and click the **Exceptions** tab to view configured exceptions. If no exceptions are configured, you can filter for threat signatures or select **Show all signatures**.

STEP 3 | Hover over a **Threat Name** or the threat **ID** to open the drop-down, and click **Exception** to review both the threat details and how the firewall is configured to enforce the threat.

For example, find out more about a top threat charted on the ACC:



STEP 4 | Review the latest **Threat Details** for the threat and launch a Threat Vault search based on the threat ID.

© 2018 Palo Alto Networks, Inc.

- Threat details displayed include the latest Threat Vault information for the threat, resources you can use to learn more about the threat, and CVEs associated with the threat.
- Select **View in Threat Vault** to open a Threat Vault search in a new window and look up the latest information the Palo Alto Networks threat database has for this threat signature.

STEP 5 | Check if a threat signature is configured as an exception to your security policy.

- If the **Used in current security rule** column is clear, the firewall is enforcing the threat based on the recommended default signature action (for example, block or alert).
- A checkmark anywhere in the **Used in current security rule** column indicates that a security policy rule is configured to enforce a non-default action for the threat (for example, allow), based on the associated **Exempt Profiles** settings.

The Used in security rule column does not indicate if the Security policy rule is enabled, only if the Security policy rule is configured with the threat exception. Select Policies > Security to check if an indicated security policy rule is enabled.

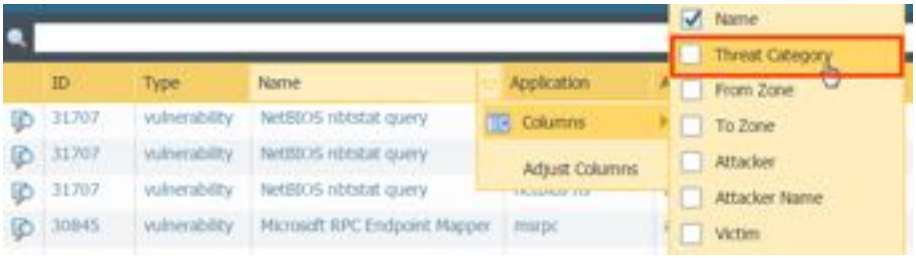
STEP 6 | **Add** an IP address on which to filter the threat exception or view existing **Exempt IP Addresses**.

Configure an exempt IP address to enforce a threat exception only when the associated session has either a matching source or destination IP address; for all other sessions, the threat is enforced based on the default signature action.

Monitor Activity and Create Custom Reports Based on Threat Categories

Threat categories classify different types of threat signatures to help you understand and draw connections between events threat signatures detect. Threat categories are subsets of the more broad threat signature types: spyware, vulnerability, antivirus, and DNS signatures. Threat log entries display the **Threat Category** for each recorded event.

- Filter Threat logs by threat category.
  1. Select **Monitor** > **Logs** > **Threat**.
  2. Add the Threat Category column so you can view the Threat Category for each log entry:

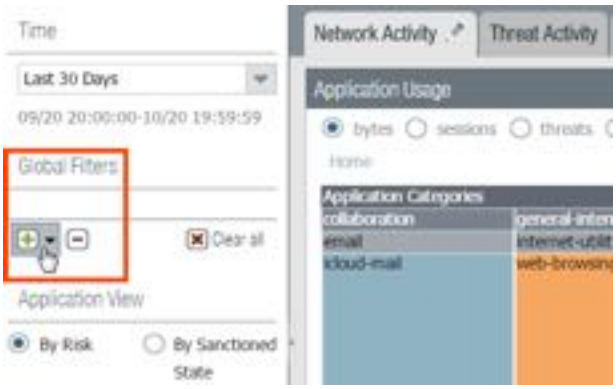


  3. To filter based on Threat Category:
     - Use the log query builder to add a filter with the **Attribute** Threat Category and in the **Value** field, enter a Threat Category.
     - Select the Threat Category of any log entry to add that category to the filter:

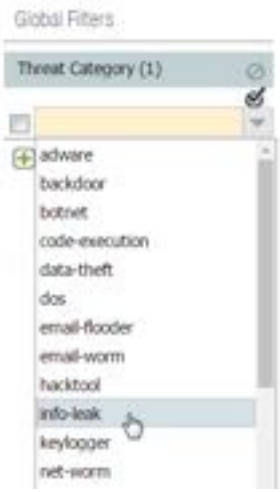
© 2018 Palo Alto Networks, Inc.



- Filter ACC activity by threat category.
  1. Select **ACC** and add Threat Category as a global filter:



  2. Select the Threat Category to filter all ACC tabs.



- Create custom reports based on threat categories to receive information about specific types of threats that the firewall has detected.
  1. Select **Monitor** > **Manage Custom** reports to add a new custom report or modify an existingone.
  2. Choose the **Database** to use as the source for the custom report—in this case, select **Threat** from either of the two types of database sources, summary databases and Detailed logs. Summary database data is condensed to allow a faster response time when generating reports. Detailed logs take longer to generate but provide an itemized and complete set of data for each log entry.
  3. In the Query Builder, add a report filter with the Attribute **Threat Category** and in the Value field, select a threat category on which to base your report.
  4. To test the new report settings, click **Run Now**.
  5. Click **OK** to save the report.

© 2018 Palo Alto Networks, Inc.

Content Delivery Network Infrastructure for Dynamic Updates

Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to the Palo Alto Networks firewalls. The firewalls access the web resources in the CDN to perform various App-ID and Content-ID functions. For enabling and scheduling the content updates, see Install Content and Software Updates.

The following table lists the web resources that the firewall accesses for a feature or application:

Resource	URL	Static Addresses (If a static server is required)
Application Database	• updates.paloaltonetworks.com:443	staticupdates.paloaltonetworks.com or the IP address 199.167.52.15
Threat/Antivirus Database	• updates.paloaltonetworks.com:443 • downloads.paloaltonetworks.com:443 As a best practice, set the update server to updates.paloaltonetworks.com. This allows the Palo Alto Networks firewall to receive content updates from the server closest to it in the CDN infrastructure.	staticupdates.paloaltonetworks.com or the IP address 199.167.52.15
PAN-DB URL Filtering	*.urlcloud.paloaltonetworks.com Resolves to the primary URL s0000.urlcloud.paloaltonetworks.com and is then redirected to the regional server that is closest: • s0100.urlcloud.paloaltonetworks.com • s0200.urlcloud.paloaltonetworks.com • s0300.urlcloud.paloaltonetworks.com • s0500.urlcloud.paloaltonetworks.com	Static IP addresses are not available. However, you can manually resolve a URL to an IP address and allow access to the regional server IP address.
BrightCloud URL Filtering	• database.brightcloud.com:443/80 • service.brightcloud.com:80	Contact BrightCloud Customer Support.

© 2018 Palo Alto Networks, Inc.

Threat Prevention Resources

For more information on Threat Prevention, refer to the following sources:

- Creating Custom Threat Signatures
- Threat Prevention Deployment
- Understanding DoS Protection

To view a list of threats and applications that Palo Alto Networks products can identify, use the following links:

- Applipedia—Provides details on the applications that Palo Alto Networks can identify.
- Threat Vault—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

© 2018 Palo Alto Networks, Inc.

Decryption

Palo Alto Networks firewalls can decrypt and inspect traffic to provide visibility into threats and to control protocols, certificate verification, and failure handling. Decryption can enforce policies on encrypted traffic so that the firewall handles encrypted traffic according to your configured security settings. Decrypt traffic to prevent malicious encrypted content from entering your network and sensitive content from leaving your network concealed as encrypted traffic. Enabling decryption can include preparing the keys and certificates required for decryption, creating decryption profiles and policies, and configuring decryption port mirroring.

- Decryption Overview
- Decryption Concepts
- Prepare to Deploy Decryption
- Define Traffic to Decrypt
- Configure SSL Forward Proxy
- Configure SSL Inbound Inspection
- Configure SSH Proxy
- Configure Server Certificate Verification for Undecrypted Traffic
- Decryption Exclusions
- Enable Users to Opt Out of SSL Decryption
- Configure Decryption Port Mirroring
- Verify Decryption
- Temporarily Disable SSL Decryption

© 2018 Palo Alto Networks, Inc.

Decryption Overview

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 and SSH encapsulate traffic, encrypting data so that it is meaningless to entities other than the client and server with the certificates to affirm trust between the devices and the keys to decode the data. Decrypt SSL and SSH traffic to:

- Prevent malware concealed as encrypted traffic from being introduced into your network. For example, an attacker compromises a website that uses SSL encryption. Employees visit that website and unknowingly download an exploit or malware. The malware then uses the infected employee endpoint to move laterally through the network and compromise other systems.
- Prevent sensitive information from moving outside the network.
- Ensure the appropriate applications are running on a secure network.
- Selectively decrypt traffic; for example, create a Decryption policy and profile to exclude traffic for financial or healthcare sites from decryption.

Palo Alto Networks firewall decryption is policy-based, and can decrypt, inspect, and control inbound and outbound SSL and SSH connections. A Decryption policy enables you to specify traffic to decrypt by destination, source, service, or URL category, and to block, restrict, or forward the specified traffic according to the security settings in the associated Decryption profile. 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. The firewall uses certificates and keys to decrypt traffic to plaintext, and then enforces App-ID and security settings on the plaintext traffic, including Decryption, Antivirus, Vulnerability, Anti-Spyware, URL Filtering, WildFire, and File-Blocking profiles. After decrypting and inspecting traffic, the firewall re-encrypts the plaintext traffic as it exits the firewall to ensure privacy and security.

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, such as 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.



Use the Decryption Best Practices Checklist to plan, implement, and maintain your decryption deployment.

You can integrate a hardware security module (HSM) with a firewall to enable enhanced security for the private keys used in SSL forward proxy and SSL inbound inspection decryption. To learn more about storing and generating keys using an HSM and integrating an HSM with your firewall, see Secure Keys with a Hardware Security Module.

You can also use Decryption Mirroring to forward decrypted traffic as plaintext to a third party solution for additional analysis and archiving.



If you enable Decryption mirroring, be aware of local laws and regulations about what traffic you can mirror and where and how you can store the traffic, because all mirrored traffic, including sensitive information, is forwarded in cleartext.

© 2018 Palo Alto Networks, Inc.

Decryption Concepts

Learn about keys and certificates for decryption, decryption policies, decryption profiles, and decryption port mirroring:

- Keys and Certificates for Decryption Policies
- SSL Forward Proxy
- SSL Forward Proxy Decryption Profile
- SSL Inbound Inspection
- SSL Inbound Inspection Decryption Profile
- SSL Protocol Settings Decryption Profile
- SSH Proxy
- SSH Proxy Decryption Profile
- Decryption Profile for No Decryption
- Decryption Mirroring
- SSL Decryption for Elliptical Curve Cryptography (ECC) Certificates
- Perfect Forward Secrecy (PFS) Support for SSL Decryption
- High Availability (HA) Support for Decrypted Sessions

Keys and Certificates for Decryption Policies

Keys are strings of numbers typically generated using a mathematical operation involving random numbers and large primes. Keys transform strings—such as passwords and shared secrets—from unencrypted plaintext to encrypted ciphertext and from encrypted ciphertext to unencrypted plaintext. Keys can be symmetric (the same key is used to encrypt and decrypt) or asymmetric (one key is used for encryption and a mathematically related key is used for decryption). Any system can generate a key.

X.509 certificates establish trust between a client and a server to establish an SSL connection. A client attempting to authenticate a server (or a server authenticating a client) knows the structure of the X.509 certificate and therefore knows how to extract identifying information about the server from fields within the certificate, such as the FQDN or IP address (called a *common name* or *CN* within the certificate) or the name of the organization, department, or user to which the certificate was issued. A certificate authority (CA) must issue all certificates. After the CA verifies a client or server, the CA issues the certificate and signs it with a private key.

When you apply a decryption policy to traffic, a session between the client and the server is established only if the firewall trusts the CA that signed the server certificate. In order 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, the firewall will present 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.

For detailed information on certificates, see Certificate Management.



To control the trusted CAs that your firewall trusts, use the Device > Certificate Management > Certificates > Default Trusted Certificate Authorities tab on the firewall web interface.

The following table describes the different certificates Palo Alto Networks firewalls use for decryption.

© 2018 Palo Alto Networks, Inc.

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. To configure a Forward Trust certificate on the firewall to present to clients when the server certificate is signed by a trusted CA, see Configure SSL Forward Proxy. 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 (see Store Private Keys on an HSM).  *Back up the private key associated with the firewall's Forward Trust CA certificate (not the firewall's master key) in a secure repository so that if an issue occurs with the firewall, you can still access the Forward Trust CA certificate. For added security, consider storing the private key associated with the Forward Trust certificate on a hardware security module (see Store Private Keys on an HSM).*
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. To configure a Forward Untrust certificate on the firewall, see Configure SSL Forward Proxy.
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.  *Beginning in PAN-OS 8.0, firewalls use the Elliptic-Curve Diffie-Hellman Ephemeral (ECDHE) algorithm to perform strict certificate checking. This means that if the firewall uses an intermediate certificate, you must reimport the certificate from your web server to the firewall after you upgrade to a PAN-OS 8.0 or later release and combine the server certificate with the intermediate certificate (install a chained certificate). Otherwise, SSL Inbound Inspection sessions that have an intermediate certificate in the chain will fail. To install a chained certificate:* 1. *Open each certificate (.cer) file in a plain-text editor such as Notepad.* 2. *Paste each certificate end-to-end with the Server Certificate at the top with each signer included below.*

© 2018 Palo Alto Networks, Inc.

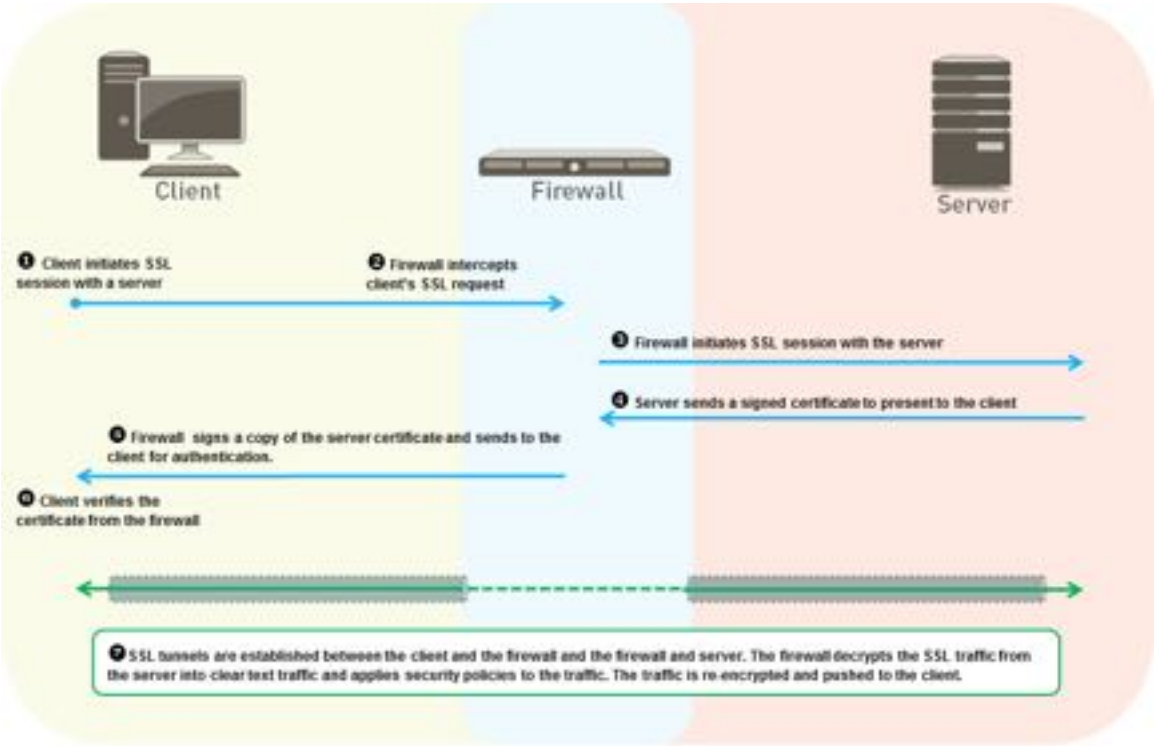
Certificates Used With Decryption	Description
	3. *Save the file as a text (.txt) or certificate (.cer) file (the name of the file cannot contain blank spaces).* 4. *Import the combined (chained) certificate into the firewall.*

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 your corporate 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 between the internal client and the external server. The firewall uses certificates to transparently represent the client to the server and to transparently represent 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 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 (for details on certificates, see Keys and Certificates for Decryption Policies).

The following figure shows this process in detail. See Configure SSL Forward Proxy for details on configuring SSL Forward Proxy.



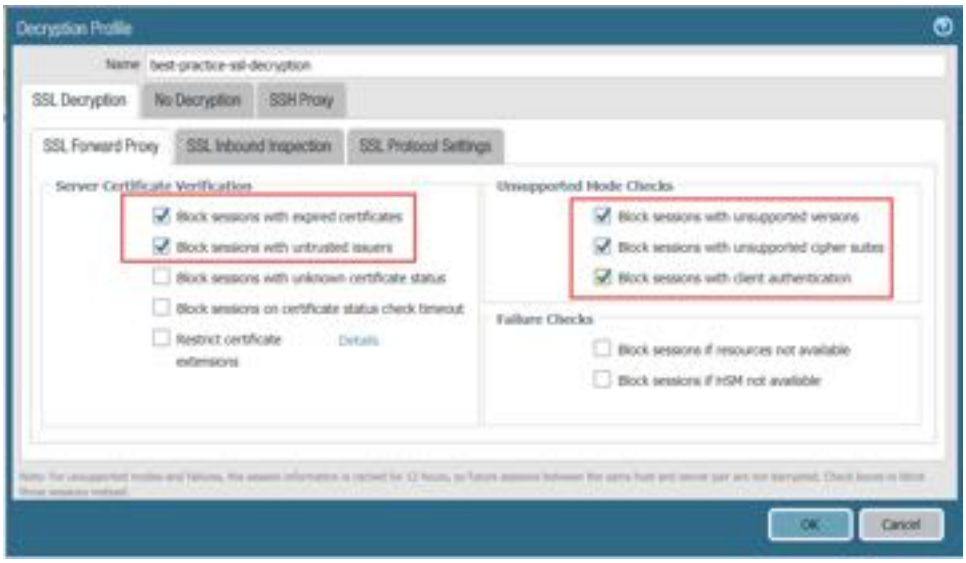
1. The internal client on your network attempts to initiate a TLS session with an external server.
2. The firewall intercepts the client's SSL certificate request. For the client, the firewall acts as the external server, even though the secure session being established is with the firewall, not with the actual server.

© 2018 Palo Alto Networks, Inc.

3. The firewall then forwards the client's SSL certificate request to the server to initiate a separate session with the server. To the server, the firewall looks like the client, the server doesn't know there's a man-in-the-middle, and the server verifies the certificate.
4. The server sends the firewall a signed certificate intended for the client.
5. The firewall analyzes the server certificate. If the server certificate is signed by a CA that the firewall trusts and meets the policies and profiles you configure, the firewall generates an SSL Forward Trust copy of the server certificate and sends it to the client. If the server certificate is signed by a CA that the firewall does not trust, the firewall generates an SSL Forward Untrust copy of the server certificate and sends it to the client. The certificate copy the firewall generates and sends to the client contains extensions from the original server certificate and is called an *impersonation* certificate because it is not the server's actual certificate. If the firewall does not trust the server, the client sees a block page warning message that the site they're attempting to connect to is not trusted, and if you Enable Users to Opt Out of SSL Decryption, the client can choose to proceed or terminate the session.
6. The client verifies the firewall's impersonation certificate. The client then initiates a session key exchange with the server, which the firewall proxies in the same manner as it proxies the certificates. The firewall forwards the client key to the server, and makes an impersonation copy of the server key for the client, so that firewall remains an "invisible" proxy, the client and server believe their session is with each other, but there are still two separate sessions, one between the client and the firewall, and the other between the firewall and the server. Now all parties have the certificates and keys required and the firewall can decrypt the traffic.
7. All SSL session traffic between goes through the firewall transparently between the client and the server. The firewall decrypts the SSL traffic, applies security policies and profiles and decryption profiles to the traffic, re-encrypts the traffic, and then forwards it.

SSL Forward Proxy Decryption Profile

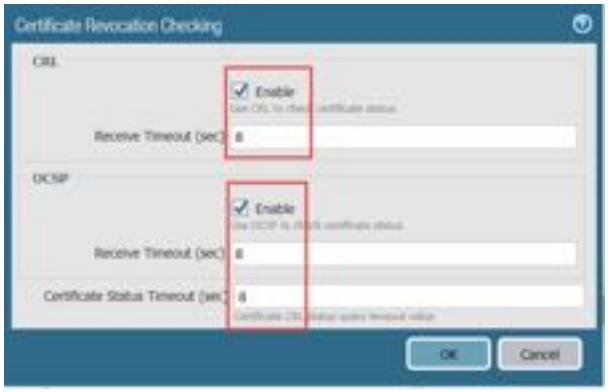
The SSL Forward Proxy Decryption profile (**Objects** > **Decryption Profile** > **SSL Decryption** > **SSL Forward Proxy**) controls the server verification, session mode checks, and failure checks for outbound traffic defined in SSL Forward Proxy Decryption policies to which you attach the profile. The following figure shows the general best practice recommendations for SSL Forward Proxy Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations. There are also specific best practices for perimeter internet gateway decryption profiles and for data center decryption profiles.



Server Certificate Verification:

© 2018 Palo Alto Networks, Inc.

- **Block sessions with expired certificates**—Always check this box to block sessions with servers that have expired certificates and prevent access to potentially insecure sites. If you don't check this box, users can connect with and transact with potentially malicious sites and see warning messages when they attempt to connect, but the connection is not prevented.
- **Block sessions with untrusted issuers**—Always check this box to block sessions with servers that have untrusted certificate issuers. An untrusted issuer may indicate a man-in-the-middle attack, a replay attack, or other attack.
- **Block sessions with unknown certificate status**—Blocks the SSL session when a the certificate revocation status of the server returns with the status "unknown". Because certificate status may be unknown for multiple reasons, for general decryption security, checking this box usually tightens security too much. However, in higher-security areas of the network such as the data center, checking this box makes sense.
- **Block sessions on certificate status check timeout**—Whether to block sessions if the status check times out depends on your company's security compliance stance because it's a tradeoff between tighter security and a better user experience. Certificate status verification examines the Certificate Revocation List (CRL) on a revocation server or uses Online Certificate Status Protocol (OCSP) to find out if the issuing CA has revoked the certificate and the certificate should not be trusted. However, revocation servers can be slow to respond, which can cause the session to timeout and the firewall to block the session even though the certificate may be valid. If you **Block sessions on certificate status check timeout** and the revocation server is slow to respond, you can use **Device** > **Setup** > **Session** > **Decryption Settings** and click **Certificate Revocation Checking** to change the default timeout value of five seconds to another value. For example, you could increase the timeout value to eight seconds, as shown in the following figure. Enable both CRL and OCSP certificate revocation checking because server certificates can contain the CRL URL in the CRL Distribution Point (CDP) extension or the OCSP URL in the Authority Information Access (AIA) certificate extension.



- **Restrict certificate extensions**—Checking this box limits the certificate extensions in the server certificate to key usage and extended key usage and blocks certificates with other extensions. However, in certain deployments, some other certificate extensions may be necessary, so only check this box if your deployment requires no other certificate extensions.

Unsupported Mode Checks. If you don't block sessions with unsupported modes, users receive a warning message if they connect with potentially unsafe servers, and they can click through that message and reach the potentially dangerous site. Blocking these sessions protects you from servers that use weak, risky protocol versions and algorithms:

- **Block sessions with unsupported versions**—When you configure the SSL Protocol Settings Decryption Profile, you specify the minimum version of SSL protocol to allow on your network to reduce the attack surface by blocking weak protocols. Always check this box to block sessions with the weak SSL protocol versions that you have chosen not to support.
- **Block sessions with unsupported cipher suites**—Always check this box to block sessions if the firewall doesn't support the cipher suite specified in the SSL handshake. You configure which algorithms the firewall supports on the **SSL Protocol Settings** tab of the Decryption profile.

© 2018 Palo Alto Networks, Inc.

- **Block sessions with client authentication**—If you have no critical applications that require client authentication, block it because firewall can't decrypt sessions that require client authentication. The firewall needs both the client and the server certificates to perform bi-directional decryption, but the firewall only knows the server certificate. This breaks decryption for client authentication sessions. When you check this box, the firewall blocks all sessions with client authentication except sessions from sites on the SSL Decryption Exclusion list (**Device** > **Certificate Management** > **SSL Decryption Exclusion**).

  If you don't **Block sessions with client authentication**, when the firewall attempts to decrypt a session that uses client authentication, the firewall allows the session and adds an entry in its local decrypt exclude cache that contains the server URL/IP address, the application, and the Decryption profile. Entries remain in the cache for 12 hours and then age out. If the same user or a different user attempts to access the serer within 12 hours using client authentication, the firewall matches the session to the decrypt exclude cache entry, does not attempt to decrypt the traffic, and allows the encrypted session.

  If the exclude cache becomes full, the firewall purges the oldest entries as new entries arrive. If you change the Decryption policy or profile, the firewall flushes the exclude cache because changing the policy or profile can change the classification outcome of the session.

  You may need to allow traffic on your network from other sites that use client authentication in addition to the Predefined sites on the SSL Decryption Exclusion list. Create a Decryption profile that allows sessions with client authentication. Add it to a Decryption policy rule that applies only to the server(s) that house the application. To increase security even more, you can require Multi-Factor Authentication to complete the user login process.

Failure Checks:

- **Block sessions if resources not available**—If you don't block sessions when firewall processing resources aren't available, then encrypted traffic that you want to decrypt enters the network still encrypted, risking allowing potentially dangerous connections. However, blocking sessions when firewall processing resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement failure checks depends on your company's security compliance stance and the importance to your business of the user experience, weighed against tighter security. Alternatively, consider using firewall models with more processing power so that you can decrypt more traffic.
- **Block sessions if HSM not available**—If you use a Hardware Security Module (HSM) to store your private keys, whether you use one depends on your compliance rules about where the private key must come from and how you want to handle encrypted traffic if the HSM isn't available. For example, if your company mandates the use of an HSM for private key signing, then block sessions if the HSM isn't available. However, if your company is less strict about this, then you can consider not blocking sessions if the HSM isn't available. (If the HSM is down, the firewall can process decryption for sites for which it has cached the response from the HSM, but not for other sites.) The best practice in this case depends on your company's policies. If the HSM is critical to your business, run the HSM in a high-availability (HA) pair (PAN-OS 8.0 supports two members in an HSM HA pair).

SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL/TLS traffic from a client to a targeted network server (any server you have the certificate for and can import it onto the firewall) and block suspicious sessions. For example, if an employee is remotely connected to a web server hosted on the company network and is attempting to add restricted internal documents to his Dropbox folder (which uses SSL for data transmission), SSL Inbound Inspection can ensure that the sensitive data does not move outside the secure company network by blocking or restricting the session.

On the firewall, you must install the certificate and private key for each server for which you want to perform SSL inbound inspection. You must also install the public key certificate as well as the private key on

© 2018 Palo Alto Networks, Inc.

each firewall that performs SSL inbound inspection. The way the firewall performs SSL inbound inspection depends on the type of key negotiated, Rivest, Shamir, Adleman (RSA) or Perfect Forward Secrecy (PFS).

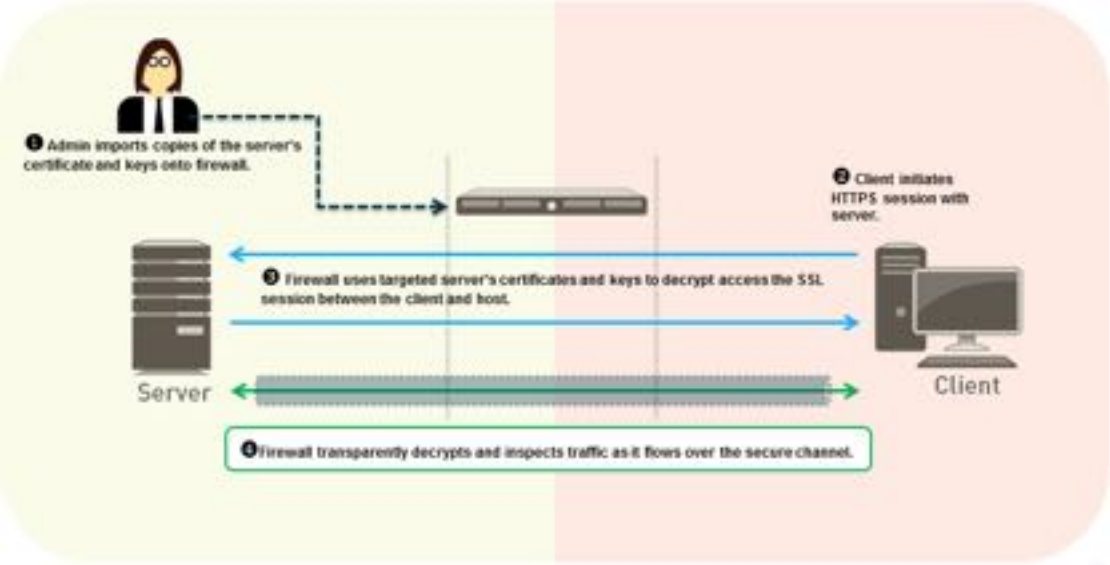
For RSA keys, the firewall performs SSL inbound inspection without terminating the connection. As the encrypted session flows through the firewall, the firewall transparently makes a copy of it and decrypts it so that the firewall can apply the appropriate policy to the traffic.

When you configure the SSL Protocol Settings Decryption Profile for SSL Inbound Inspection traffic, create separate profiles for servers with different security capabilities. For example, if one set of servers supports only RSA, the SSL Protocol Settings only need to support RSA. However, the SSL Protocol Settings for servers that support PFS should support PFS. Configure SSL Protocol Settings for the highest level of security that the server supports, but check performance to ensure that the firewall resources can handle the higher processing load that higher security protocols and algorithms require.

For PFS keys using the Diffie-Hellman exchange (DHE) or Elliptic Curve Diffie-Hellman exchange (ECDHE), the firewall acts as a man-in-the-middle proxy between the external client and the internal server. Because PFS generates a new key with every session, the firewall can't simply copy and decrypt the inbound SSL flow as it passes through, the firewall must act as a proxy device.

The following figure shows how SSL Inbound Inspection works when the key exchange algorithm is RSA. When the key exchange algorithm is PFS, the firewall functions as a proxy (creates a secure session between the client and the firewall and another secure session between the firewall and the server) and must generate a new session key for each secure session.

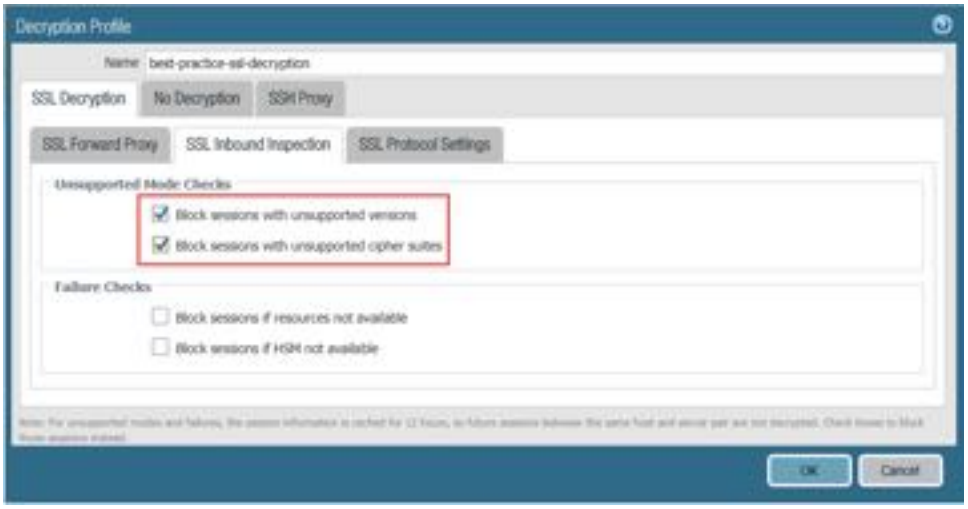
See Configure SSL Inbound Inspection for details on enabling this feature.



SSL Inbound Inspection Decryption Profile

The SSL Inbound Inspection Decryption profile (**Objects** > **Decryption Profile** > **SSL Decryption** > **SSL Inbound Inspection**) controls the session mode checks and failure checks for inbound traffic defined in the SSL Inbound Inspection Decryption policies to which you attach the profile. The following figure shows the general best practice recommendations for SSL Inbound Inspection Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations.

© 2018 Palo Alto Networks, Inc.



Unsupported Mode Checks. If you don't block sessions with unsupported modes, users receive a warning message if they connect with potentially unsafe servers, and they can click through that message and reach the potentially dangerous site. Blocking these sessions protects you from servers that use weak, risky protocol versions and algorithms:

1. **Block sessions with unsupported versions**—When you configure the SSL Protocol Settings Decryption Profile, you specify the minimum version of SSL protocol to allow on your network to reduce the attack surface by blocking weak protocols. Always check this box to block sessions with the weak SSL protocol versions that you have chosen not to support.
2. **Block sessions with unsupported cipher suites**—Always check this box to block sessions if the firewall doesn't support the cipher suite specified in the SSL handshake. You configure which algorithms the firewall supports on the **SSL Protocol Settings** tab of the Decryption profile.

Failure Checks:

- **Block sessions if resources not available**—If you don't block sessions when firewall processing resources aren't available, then encrypted traffic that you want to decrypt enters the network still encrypted, risking allowing potentially dangerous connections. However, blocking sessions when firewall processing resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement failure checks depends on your company's security compliance stance and the importance to your business of the user experience, weighed against tighter security. Alternatively, consider using firewall models with more processing power so that you can decrypt more traffic.
- **Block sessions if HSM not available**—If you use a Hardware Security Module (HSM) to store your private keys, whether you use one depends on your compliance rules about where the private key must come from and how you want to handle encrypted traffic if the HSM isn't available. For example, if your company mandates the use of an HSM for private key signing, then block sessions if the HSM isn't available. However, if your company is less strict about this, then you can consider not blocking sessions if the HSM isn't available. (If the HSM is down, the firewall can process decryption for sites for which it has cached the response from the HSM, but not for other sites.) The best practice in this case depends on your company's policies. If the HSM is critical to your business, run the HSM in a high-availability (HA) pair (PAN-OS 8.0 supports two members in an HSM HA pair).

SSL Protocol Settings Decryption Profile

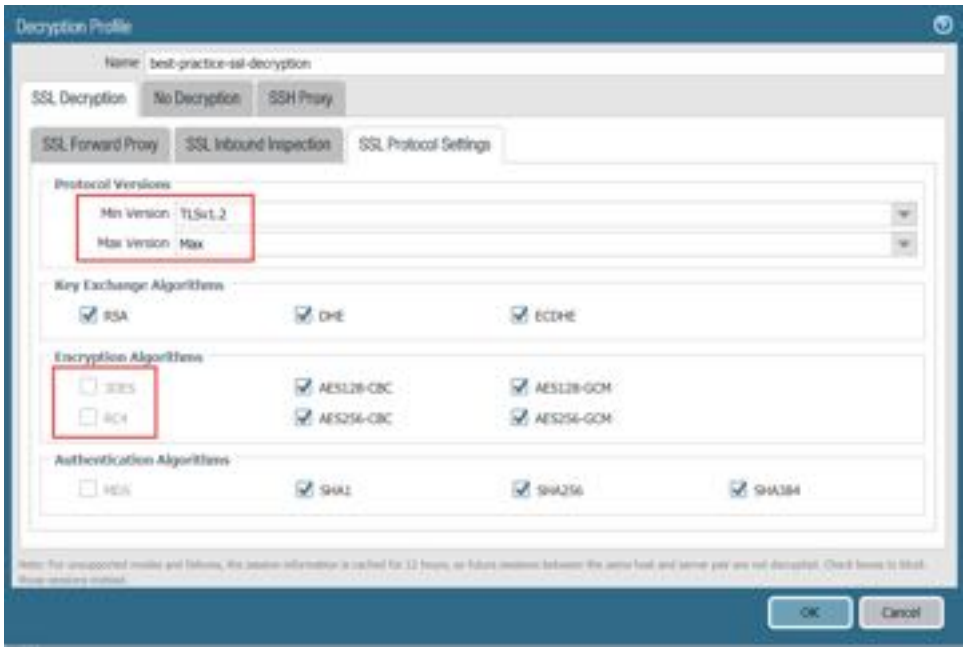
The SSL Protocol Settings (**Objects** > **Decryption Profile** > **SSL Decryption** > **SSL Protocol Settings**) control whether you allow vulnerable SSL/TLS protocol versions, weak encryption algorithms, and weak authentication algorithms. SSL Protocol Settings apply to outbound SSL Forward Proxy and inbound SSL

© 2018 Palo Alto Networks, Inc.

Inbound Inspection traffic. These settings don't apply to SSH Proxy traffic or to traffic that you don't decrypt.

The following figure shows the general best practice recommendations for SSL Protocol Settings. There are also specific best practices for perimeter internet gateway decryption profiles and for data center decryption profiles.

When you configure SSL Protocol Settings for SSL Inbound Inspection traffic, create separate profiles for servers with different security capabilities. For example, if one set of servers supports only RSA, the SSL Protocol Settings only need to support RSA. However, the SSL Protocol Settings for servers that support PFS should support PFS. Configure SSL Protocol Settings for the highest level of security that the target server you are protecting supports, but check performance to ensure that the firewall resources can handle the higher processing load that higher security protocols and algorithms require.



Protocol Versions:

- Set the **Min Version** to **TLSv1.2** to provide the strongest security—business sites that value security support TLSv1.2. If a site (or a category of sites) only supports weaker ciphers, review the site and determine if it really houses a legitimate business application. If it does, make an exception for only that site by configuring a Decryption profile with a **Min Version** that matches the strongest cipher the site supports and then applying the profile to a Decryption policy rule that limits allowing the weak cipher to only the site or sites in question. If the site doesn't house a legitimate business application, don't weaken your security posture to support the site—weak protocols (and ciphers) contain known vulnerabilities that attackers can exploit. If the site belongs to a category of sites that you don't need for business purposes, use URL Filtering to block access to the entire category. Don't support weak encryption or authentication algorithms unless you must do so to support important legacy sites, and when you make exceptions, create a separate Decryption profile that allows the weaker protocol just for those sites. Don't downgrade the main Decryption profile that you apply to most sites to TLSv1.1 just to accommodate a few exceptions.

  Decrypting TLS traffic forces browsers that use HTTP/2 to fall back to HTTP 1.1 because the firewall can't decrypt HTTP/2 traffic. Allow browsers to fall back to HTTP 1.1 so you can decrypt this traffic and prevent potentially dangerous traffic from entering the network as encrypted traffic.

© 2018 Palo Alto Networks, Inc.



Qualys SSL Labs SSL Pulse web page provides up-to-date statistics on the percentages of different ciphers and protocols in use on the 150,000 most popular sites in the world so you can see trends and understand how widespread worldwide support is for more secure ciphers and protocols.

- Set the **Max Version** to **Max** rather than to a particular version so that as the protocols improve, the firewall automatically supports the newest and best protocols. Whether you intend to attach a Decryption profile to a Decryption policy rule that governs inbound (SSL Inbound Inspection) or outbound (SSL Forward Proxy) traffic, avoid allowing weak algorithms.

Key Exchange Algorithms: Leave all three boxes checked (default) to support both RSA and PFS (DHE and ECDHE) key exchanges.

Encryption Algorithms: When you set the protocol version to TLSv1.2, the older, weaker 3DES and RC4 algorithms are automatically unchecked (blocked). For any traffic for which you must allow a weaker TLS protocol, create a separate Decryption profile and apply it only to traffic for that site, and uncheck the 3DES and RC4 boxes. Do not allow traffic that uses the 3DES or RC4 algorithms. If unchecking the 3DES or RC4 boxes prevents you from accessing a site that you must use for business, create a separate Decryption profile for that site. Don't weaken decryption for any other sites.

Authentication Algorithms: The older, weaker MD5 algorithm is automatically unchecked (blocked). Do not allow MD5 authenticated traffic on your network; SHA1 is the weakest authentication algorithm you should allow. If no necessary sites use SHA1, uncheck the box and block traffic to further reduce the attack surface.

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 boot up process, the firewall checks if there is an existing key. If not, the firewall generates a key. The firewall uses the key to decrypt SSH sessions for all virtual systems configured on the firewall and all 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 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 packets and TCP packets. One SSH connection may 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:

- session
- X11
- forwarded-tcpip
- direct-tcpip

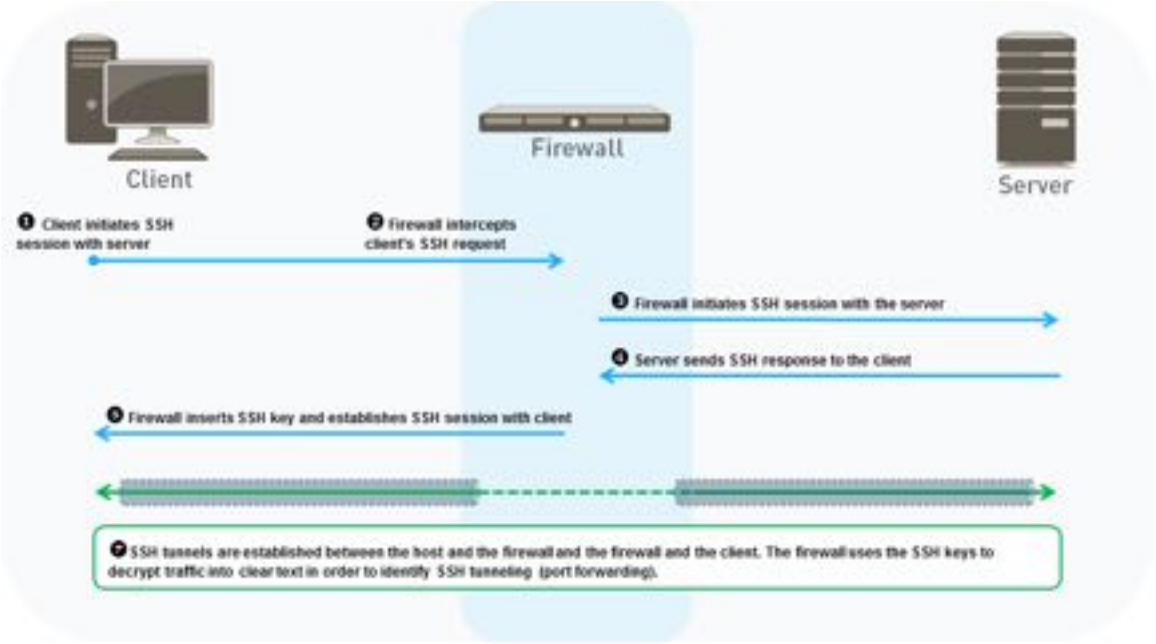
When the channel type is session, the firewall identifies the traffic as allowed SSH traffic such as SFTP or SCP. When the channel type is X11, forwarded-tcpip, or direct-tcpip, the firewall identifies the traffic as SSH tunneling traffic and blocks it.



Limit SSH use to administrators who need to manage network devices, log all SSH traffic, and consider configuring Multi-Factor Authentication to help ensure that only legitimate users can use SSH to access devices, which reduces the attack surface.

© 2018 Palo Alto Networks, Inc.

The following figure shows how SSH Proxy decryption works. See Configure SSH Proxy for details on how to enable SSH Proxy decryption.



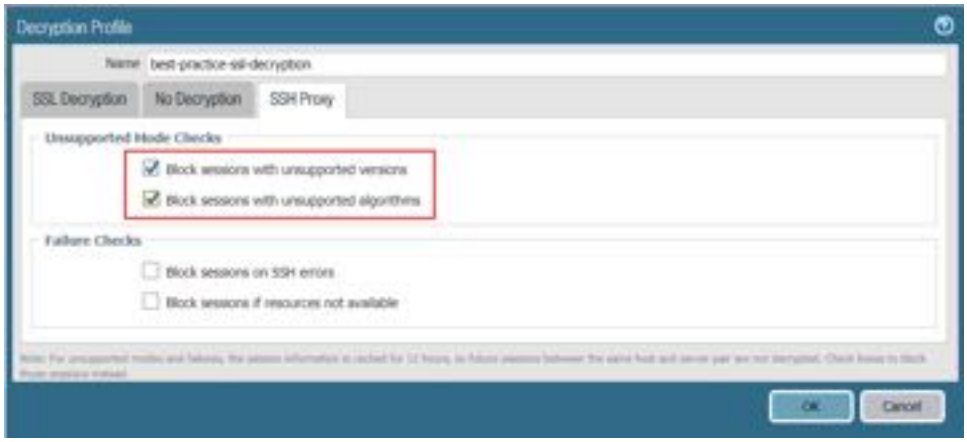
When the client sends an SSH request to the server to initiate a session, the firewall intercepts the request and forwards it to the server. The firewall then intercepts the server response and forwards it to the client. This establishes two separate SSH tunnels, one between the firewall and the client and one between the firewall and the server, with firewall functioning as a proxy. As traffic flows between the client and the server, the firewall checks whether the SSH traffic is being routed normally or if it is using SSH tunneling (port forwarding). The firewall doesn't perform content and threat inspection on SSH tunnels; however, if the firewall identifies SSH tunnels, it blocks the SSH tunneled traffic and restricts the traffic according to configured security policies.

SSH Proxy Decryption Profile

The SSH Proxy Decryption profile (**Objects** > **Decryption Profile** > **SSH Proxy**) controls the session mode checks and failure checks for 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 SSH Proxy Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations.

The firewall doesn't perform content and threat inspection on SSH tunnels (port forwarding). However, the firewall distinguishes between the SSH application and the SSH-tunnel application. If the firewall identifies SSH tunnels, it blocks the SSH tunneled traffic and restricts the traffic according to configured security policies.

© 2018 Palo Alto Networks, Inc.



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, and they can click through that message and reach the potentially dangerous site. Blocking these sessions protects you from servers that use weak, risky protocol versions and algorithms:

1. **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.
2. **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.

Failure Checks:

- **Block sessions on SSH errors**—Checking this box terminates the session if SSH errors occur.
- **Block sessions if resources not available**—If you don't block sessions when firewall processing resources aren't available, then encrypted traffic that you want to decrypt enters the network still encrypted, risking allowing potentially dangerous connections. However, blocking sessions when firewall processing resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement failure checks depends on your company's security compliance stance and the importance to your business of the user experience, weighed against tighter security. Alternatively, consider using firewall models with more processing power so that you can decrypt more traffic.

Decryption Profile for No Decryption

The No Decryption profile (**Objects** > **Decryption Profile** > **No Decryption**) controls server verification checks for traffic that you choose not to decrypt as defined in "No Decryption" Decryption policies to which you attach the profile. (Don't exclude traffic that you can't decrypt because a site breaks decryption for technical reasons such as a pinned certificate or mutual authentication by policy. Instead, add the hostname to the Decryption Exclusion List.) The following figure shows the general best practice recommendations for the No Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations.

© 2018 Palo Alto Networks, Inc.



- **Block sessions with expired certificates**—Always check this box to block sessions with servers that have expired certificates and prevent access to potentially insecure sites. If you don't check this box, users can connect with and transact with potentially malicious sites and see warning messages when they attempt to connect, but the connection is not prevented.
- **Block sessions with untrusted issuers**—Always check this box to block sessions with servers that have untrusted certificate issuers. An untrusted issuer may indicate a man-in-the-middle attack, a replay attack, or other attack.

Decryption Mirroring

The decryption mirroring feature provides the capability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as NetWitness or Solera—for archiving and analysis. This feature is necessary for organizations that require comprehensive data capture for forensic and historical purposes or data leak prevention (DLP) functionality. Decryption mirroring is available on PA-7000 Series, PA-5200 Series, PA-5000 Series and PA-3000 Series platforms only and requires that a free license be installed to enable this feature.

Keep in mind that the decryption, storage, inspection, and/or use of SSL traffic is regulated in certain countries and user consent may be required in order to use the decryption mirror feature. Additionally, use of this feature could enable malicious users with administrative access to the firewall to harvest usernames, passwords, social security numbers, credit card numbers, or other sensitive information submitted using an encrypted channel. Palo Alto Networks recommends that you consult with your corporate counsel before activating and using this feature in a production environment.

The following graphic shows the process for mirroring decrypted traffic and the section Configure Decryption Port Mirroring describes how to license and enable this feature.



SSL Decryption for Elliptical Curve Cryptography (ECC) Certificates

The firewall automatically decrypts SSL traffic from websites and applications using ECC certificates, including Elliptical Curve Digital Signature Algorithm (ECDSA) certificates. As organizations transition to using ECC certificates to benefit from the strong keys and small certificate size, you can continue to maintain visibility into and safely enable ECC-secured application and website traffic.

© 2018 Palo Alto Networks, Inc.

Decryption for websites and applications using ECC certificates is not supported for traffic that is mirrored to the firewall; encrypted traffic using ECC certificates must pass through the firewall directly for the firewall to decrypt it.

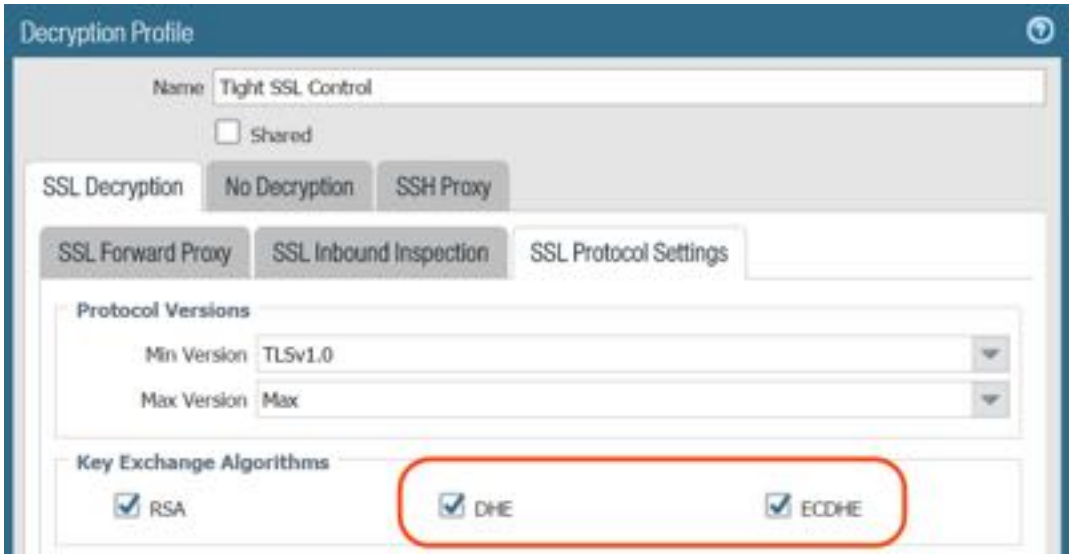
You cannot use a hardware security module (HSM) to store the private keys associated with ECDSA certificates.

Perfect Forward Secrecy (PFS) Support for SSL Decryption

PFS is a secure communication protocol that prevents the compromise of one encrypted session from leading to the compromise of multiple encrypted sessions. With PFS, a server generates unique private keys for each secure session it establishes with a client. If a server private key is compromised, only the single session established with that key is vulnerable—an attacker cannot retrieve data from past and future sessions because the server establishes each connection with a uniquely generated key. The firewall decrypts SSL sessions established with PFS key exchange algorithms.

Support for Diffie-Hellman (DHE)-based PFS and elliptical curve Diffie-Hellman (ECDHE)-based PFS is enabled by default (**Objects** > **Decryption Profile** > **SSL Decryption** > **SSL Protocol Settings**).

If you use the DHE or ECDHE key exchange algorithms to enable Perfect Forward Secrecy (PFS) Support for SSL Decryption, you cannot use a hardware security module(HSM) to store the private keys for SSL Inbound Inspection.



High Availability (HA) Support for Decrypted Sessions

High Availability (HA) syncs are supported for inbound, decrypted SSL sessions, if the sessions were established using non-PFS key exchange algoriothms. When a failover occurs, the passive device continues to inspect and enforce the decrypted traffic.

HA syncs are not supported for:

- decrypted SSL sessions (both inbound and outbound) that were established using PFS key exchange algorithms
- decrypted, outbound SSL sessions using non-PFS key exchange algorithms

In these cases, when a failover occurs, the passive device allows transferred sessions without decrypting them. New sessions will then continue to be decrypted based on your decryption policy.

The following table details HA support for decrypted sessions:

© 2018 Palo Alto Networks, Inc.

	PFS Protected Session	Non-PFS Protected Session
Inbound SSL Session (Inbound Inspection Decryption)	– No HA Sync	✓ HA Sync
Outbound SSL Sessions (SSL Forward Proxy Decryption)	– No HA Sync	– No HA Sync

© 2018 Palo Alto Networks, Inc.

Prepare to Deploy Decryption

The most time-consuming part of deploying decryption isn't configuring the decryption policies and profiles, it's preparing for the deployment by working with stakeholders to decide what traffic to decrypt and not to decrypt, educating your user population about changes to website access, developing a private key infrastructure (PKI) strategy, and planning a staged, prioritized rollout.

Set goals for decryption and review Decryption planning best practices checklist to ensure that you understand the recommended best practices. The best practice goal is to decrypt as much traffic as your firewall resources permit and decrypt the most important traffic first.



Migrate from port-based to application-based Security policy rules before you create and deploy Decryption policy rules. If you create Decryption rules based on port-based Security policy and then migrate to application-based Security policy, the change could cause the Decryption rules to block traffic that you intend to allow because Security policy rules are likely to use application default ports to prevent application traffic from using non-standard ports. For example, traffic identified as web-browsing application traffic (default port 80) may have underlying applications that have different default ports, such as HTTPS traffic (default port 443). The application-default rule blocks the HTTPS traffic because it sees the decrypted traffic using a "non-standard" port (443 instead of 80). Migrating to App-ID based rules before deploying decryption means that when you test your decryption deployment in POCs, you'll discover Security policy misconfiguration and fix it before rolling it out to the general user population.

To prepare to deploy Decryption:

- Work with Stakeholders to Develop a Decryption Deployment Strategy
- Develop a PKI Rollout Plan
- Size the Firewall Decryption Deployment
- Plan a Staged, Prioritized Deployment

Work with Stakeholders to Develop a Decryption Deployment Strategy

Work with stakeholders such as legal, finance, HR, executives, security, and IT/support to develop a decryption deployment strategy. Start by getting the required approvals to decrypt traffic to secure the corporation. Decrypting traffic involves understanding how legal regulations and business needs affect what you can and can't decrypt.

Identify and prioritize the traffic you want to decrypt. The best practice is to decrypt as much traffic as you can to gain visibility into potential threats in encrypted traffic and prevent those threats. If incorrect firewall sizing prevents you from decrypting all of the traffic you want to decrypt, prioritize the most critical servers, the highest-risk traffic categories, and less trusted segments and IP subnets. To help prioritize, ask yourself questions such as, "What happens if this server is compromised?" and "How much risk am I willing to take in relation to the level of performance I want to achieve?"

Next, identify traffic that you can't decrypt because the traffic breaks decryption for technical reasons such as a pinned certificate, unsupported ciphers, or mutual authentication. Evaluate the websites that break decryption technically and ask yourself if you need access to those sites for business reasons. If you don't need access to those sites, allow decryption to block them. If you need access to any of those sites for business purposes, add them to the SSL Decryption Exclusion list to except them from decryption. The SSL Decryption Exclusion list is exclusively for sites that break decryption technically.

© 2018 Palo Alto Networks, Inc.

Identify sensitive traffic that you *choose* not to decrypt for legal, regulatory, personal, or other reasons, such as financial, health, military, or government traffic, or the traffic of certain executives. This is not traffic that breaks decryption technically, so you don't use the SSL Decryption Exclusion list to except this traffic from decryption. Instead, you Create a Policy-Based Decryption Exclusion to identify and control traffic you choose not to decrypt and apply the No Decryption decryption profile to the policy to prevent servers with certificate issues from accessing the network. Policy-based decryption exclusions are only for traffic you choose not to decrypt.

When you plan decryption policy, consider your company's security compliance rules, computer usage policy, and your business goals. Extremely strict controls can impact the user experience by preventing access to non-business sites the user used to access, but may be required for government or financial institutions. There is always a tradeoff between usability, management overhead, and security. The tighter the decryption policy, the greater the chance that a website will become unreachable, which may result in user complaints and possibly modifying the rulebase.



Although a tight decryption policy may initially cause a few user complaints, those complaints can draw your attention to unsanctioned or undesirable websites that are blocked because they use weak algorithms or have certificate issues. Use complaints as a tool to better understand the traffic on your network.

Different groups of users and even individual users may require different decryption policies, or you may want to apply the same decryption policy to all users. For example, executives may be exempted from decryption policies that apply to other employees. And you may want to apply different decryption policies to employee groups, contracts, partners, and guests. Prepare updated legal and HR computer usage policies to distribute to all employees, contractors, partners, guests, and any other network users so that when you roll out decryption, users understand their data can be decrypted and scanned for threats.



How you handle guest users depends on the access they require. Isolate guests from the rest of your network by placing them on a separate VLAN and on a separate SSID for wireless access. If guests don't need to access your corporate network, don't let them on it and there will be no need to decrypt their traffic. If guests need to access your corporate network, decrypt their traffic:

- *Enterprises don't control guest devices. Decrypt guest traffic and subject it to your guest Security policy so the firewall can inspect the traffic and prevent threats. To do this, redirect guest users through a captive portal, instruct them how to download and install the CA certificate, and clearly notify guests that their traffic will be decrypted. Include the process in your company's privacy and computer usage policy.*
- *Create separate Decryption policy rules and Security policy rules to tightly control guest access so that guests can only access the areas of your network that they need to access.*

Similarly to different groups of users, decide which devices to decrypt and which applications to decrypt. Today's networks support not only corporate devices, but BYOD, mobile, remote-user and other devices, including contractor, partner, and guest devices. Today's users attempt to access many sites, both sanctioned and unsanctioned, and you should decide how much of that traffic you want to decrypt.



Enterprises don't control BYOD devices. If you allow BYOD devices on your network, decrypt their traffic and subject it to the same Security policy that you apply to other network traffic so the firewall can inspect the traffic and prevent threats. To do this, redirect BYOD users through a captive portal, instruct them how to download and install the CA certificate, and clearly notify users that their traffic will be decrypted. Educate BYOD users about the process and include it in your company's privacy and computer usage policy.

© 2018 Palo Alto Networks, Inc.

Decide what traffic you want to log and investigate what traffic you can log. Be aware of local laws regarding what types of data you can log and store, and where you can log and store the data. For example, local laws may prevent logging and storing personal information such as health and financial data.

Decide how to handle bad certificates. For example, will you block or allow sessions for which the certificate status is unknown? Understanding how you want to handle bad certificates determines how you configure the decryption profiles that you attach to decryption policies to control which sessions you allow based on the server certificate verification status.

Develop a PKI Rollout Plan

Plan how to roll out your public key infrastructure (PKI). Network devices need an SSL Forward Trust CA certificate for trusted sites and an SSL Forward Untrust CA certificate for untrusted sites. Generate separate Forward Trust and Forward Untrust certificates (do not sign the Forward Untrust certificate with the Enterprise Root CA because you want the Untrust certificate to warn users that they are trying to access potentially unsafe sites). Palo Alto Networks next-generation firewalls have two methods of generating CA certificates for SSL decryption:

- **Generate the SSL CA certificates from your Enterprise Root CA as subordinate certificates**—If you have an existing Enterprise PKI, this is the best practice. Generating a subordinate certificate from your Enterprise Root CA makes the rollout easier and smoother because network devices already trust the Enterprise Root CA, so you avoid any certificate issues when you begin the deployment phase. If you don't have an Enterprise Root CA, consider getting one.
- **Generate a self-signed Root CA certificate on the firewall and create subordinate CA certificates on that firewall**—If you don't have an Enterprise Root CA, this method provides a self-signed Root CA certificate and the subordinate Forward Trust and Untrust CA certificates. With this method, you need to install the self-signed certificates on all of your network devices so that those devices recognize the firewall's self-signed certificates. Because the certificates must be deployed to all devices, this method is better for small deployments and proof-of-concept (POC) trials than for large deployments.



Do not export the Forward Untrust certificate to the Certificate Trust Lists of your network devices! This is critical because installing the Untrust certificate in the Trust List results in devices trusting websites that the firewall does not trust. In addition, users won't see certificate warnings for untrusted sites, so they won't know the sites are untrusted and may access those sites, which could expose your network to threats.



Regardless of whether you generate Forward Trust certificates from your Enterprise Root CA or use a self-signed certificate generated on the firewall, generate a separate subordinate Forward Trust CA certificate for each firewall. The flexibility of using separate subordinate CAs enables you to revoke one certificate when you decommission a device (or device pair) without affecting the rest of the deployment and reduces the impact in any situation in which you need to revoke a certificate. Separate Forward Trust CAs on each firewall also helps troubleshoot issues because the CA error message the user sees includes information about the firewall the traffic is traversing. If you use the same Forward Trust CA on every firewall, you lose the granularity of that information.

There is no benefit to using different Forward Untrust certificates on different firewalls, so you can use the same Forward Untrust certificate on all firewalls. If you need additional security for your private keys, consider storing them on an HSM.

You may need to make special accommodations for guest users. If guest users don't need access to your corporate network, don't allow access, and then you won't have to decrypt their traffic or create the certificate infrastructure to support their access. If you need to allow guest users access to your corporate network, purchase a public Root CA from an organization such as Google Domains or GoDaddy so you can sign guest client certificates with it (guests are unlikely to have your Enterprise Root CA in their Certificate

© 2018 Palo Alto Networks, Inc.

Trust List so they won't trust a certificate signed by the Enterprise CA, but they will have public Root CAs in their Certificate Trust List), and decrypt guest traffic—don't allow undecrypted guest traffic on your network. Use a separate, dedicated gateway firewall for guests and use the public Root CA only on that firewall. All employees, contractors, partners, and other users who have your Enterprise Root CA in their Certificate Trust List should use the regular corporate gateway firewalls, not the firewall dedicated to guest users.

Size the Firewall Decryption Deployment

Decrypting encrypted traffic consumes firewall CPU resources and can affect throughput. In general, the tighter the security (the more SSL traffic you decrypt combined with the more stringent your protocol settings), the more firewall resources decryption consumes. Work with your Palo Alto Networks SE/CE to size your firewall deployment and avoid sizing mistakes. Factors that affect decryption resource consumption and therefore how much traffic the firewall can decrypt include:

- The amount of SSL traffic you want to decrypt. This varies from network to network. For example, some applications must be decrypted to prevent the injection of malware or exploits into the network or unauthorized data transfers, some applications can't be decrypted due to local laws and regulations or business reasons, and other applications are cleartext (unencrypted) and don't need to be decrypted. The more traffic you want to decrypt, the more resources you need.
- The TLS protocol version. Higher versions are more secure but consume more resources. Use the highest TLS protocol version you can to maximize security.
- The key size. The larger the key size, the better the security, but also the more resources the key processing consumes.
- The key exchange algorithm. Perfect Forward Secrecy (PFS) ephemeral key exchange algorithms such as Diffie-Hellman Ephemeral (DHE) Elliptic-Curve Diffie-Hellman Exchange (ECDHE) consume more processing resources than Rivest-Shamir-Adleman (RSA) algorithms. PFS key exchange algorithms provide greater security than RSA key exchange algorithms because the firewall has to generate a new cipher key for each session—but generating the new key consumes more firewall resources. However, if an attacker compromises a session key, PFS prevents the attacker from using it to decrypt any other sessions between the same client and server and RSA does not.
- The encryption algorithm. The key exchange algorithm determines whether the encryption algorithm is PFS or RSA.
- The certificate authentication method. RSA (not the RSA key exchange algorithm) consumes less resources than Elliptic Curve Digital Signature Algorithm (ECDSA) but ECDSA is more secure.



The combination of the key exchange algorithm and the certificate authentication method affect throughput performance as shown in RSA and ECDSA benchmark tests*. The performance cost of PFS trades off against the higher security that PFS achieves, but PFS may not be needed for all types of traffic. You can save firewall CPU cycles by using RSA for traffic that you want to decrypt and inspect for threats but that isn't sensitive.*

- Average transaction sizes. For example, small average transaction sizes consume more processing power to decrypt. Measure the average transaction size of all traffic, then measure the average transaction size of traffic on port 443 (the default port for HTTPS encrypted traffic) to understand the proportion of encrypted traffic going to the firewall in relation to your total traffic and the average transaction sizes. Eliminate anomalous outliers such as unusually large transactions to get a truer measurement of average transaction size.
- The firewall model and resources. Newer firewall models have more processing power than older models.

The combination of these factors determines how decryption consumes firewall processing resources. To best utilize the firewall's resources, understand the risks of the data you're protecting. If firewall resources are an issue, use stronger decryption for higher-priority traffic and use less processor-intensive decryption to decrypt and inspect lower-priority traffic until you can increase the available resources. For example,

© 2018 Palo Alto Networks, Inc.

you could use RSA instead of ECDHE and ECDSA for traffic that isn't sensitive or high-priority to preserve firewall resources for using PFS-based decryption for higher priority, sensitive traffic. (You're still decrypting and inspecting the lower-priority traffic, but trading off consuming fewer computational resources with using algorithms that aren't as secure as PFS.) The key is to understand the risks of different traffic types and treat them accordingly.

Measure firewall performance so that you understand the currently available resources, which helps you understand whether you need more firewall resources to decrypt the traffic you want to decrypt. Measuring firewall performance also sets a baseline for performance comparisons after deploying decryption.

When you size the firewall deployment, base it not only on your current needs, but also on your future needs. Include headroom for the growth of decryption traffic because Gartner predicts that through 2019, more than 80 percent of enterprise web traffic will be encrypted, and more than 50 percent of new malware campaigns will use various forms of encryption. Work with your Palo Alto Networks representatives and take advantage of their experience in sizing firewalls to help you size your firewall decryption deployment.

Plan a Staged, Prioritized Deployment

Plan to roll out decryption in a controlled manner, piece by piece. Don't roll out your entire decryption deployment at one time. Test and ensure that decryption is working as planned and that users understand what you are doing and why. Rolling out decryption in this manner makes it easier to troubleshoot in case anything doesn't work as expected and helps users adjust to the changes.

Educating stakeholders, employees, and other users such as contractors and partners is critical because decryption settings may change their ability to access some websites. Users should understand how to respond to situations in which previously reachable websites become unreachable and what information to give technical support. Support should understand what is being rolled out when and how to help users who encounter issues. Before you roll out decryption to the general population:

- Identify early adopters to help champion decryption and who will be able to help other employees who have questions during the full rollout. Enlist the help of department managers and help them understand the benefits of decrypting traffic.
- Set up proof-of-concept (POC) trials in each department with early adopters and other employees who understand why decrypting traffic is important. Educate POC participants about the changes and how to contact technical support if they run into issues. In this way, decryption POCs become an opportunity to work with technical support to POC how to support decryption and to develop the most painless method for supporting the general rollout. The interaction between POC users and technical support also allows you to fine-tune policies and how to communicate with users.

  POCs enable you to experiment with prioritizing what to decrypt first, so that when you phase in decryption in the general population, your POC experience helps you understand how to phase in decrypting different URL Categories. Measure the way decryption affects firewall CPU and memory utilization to help understand if the firewall sizing is correct or if you need to upgrade. POCs can also reveal applications that break decryption technically and need to be added to the Decryption Exclusion list.

  When you set up POCs, also set up a user group that can certify the operational readiness and procedures prior to the general rollout.
- Educate the user population before the general rollout, and plan to educate new users as they join the company. This is a critical phase of deploying decryption because the deployment may affect websites that users previously visited but are not safe, so those sites are no longer reachable. The POC experience helps identify the most important points to communicate.
- Phase in decryption. You can accomplish this several ways. You can decrypt the highest priority traffic first (for example, the URL Categories most likely to harbor malicious traffic, such as gaming) and then decrypt more as you gain experience. Alternatively, you can take a more conservative approach and

© 2018 Palo Alto Networks, Inc.

decrypt the URL Categories that don't affect your business first (so if something goes wrong, no issues occur that affect business), for example, news feeds. In all cases, the best way to phase in decryption is to decrypt a few URL Categories, take user feedback into account, run reports to ensure that decryption is working as expected, and then gradually decrypt a few more URL Categories and verify, and so on. Plan to make Decryption Exclusions to exclude sites from decryption if you can't decrypt them for technical reasons or because you choose not to decrypt them.

If you Enable Users to Opt Out of SSL Decryption (users see a response page that allows them either to opt out of decryption and end the session without going to the site or to proceed to the site and agree to have the traffic decrypted), educate them about what it is, why they're seeing it, and what their options are.

- Create realistic deployment schedules that allow time to evaluate each stage of the rollout.



Place firewalls in positions where they can see all of the network traffic so that no encrypted traffic inadvertently gains access to your network because it bypasses the firewall.

© 2018 Palo Alto Networks, Inc.

Define Traffic to Decrypt

A Decryption policy rule allows you to define traffic that you want the firewall to decrypt and to define traffic that you choose to exclude from decryption because the traffic is personal or because of local regulations, for example.

Attach a Decryption profile to each Decryption policy rule to enable certificate checks, session mode checks, failure checks, and protocol and algorithm checks, depending on the profile. These checks prevent risky connections, such as sessions with untrusted certificate issuers, weak protocols, ciphers, and algorithms, and servers that have certificate issues.



Review the Decryption deployment best practices checklist *to ensure that you understand the recommended best practices.*

As a best practice, you should always block some URL Filtering categories such as malware, phishing, dynamic-dns, unknown, command-and-control, proxy-avoidance-and-anonymizers, questionable, and parked. Many companies also block the copyright-infringement and extremism URL categories. If you must allow any of these categories for business reasons, you must decrypt them and apply strict Security profiles to the traffic.

URL categories that you should always decrypt if you allow them include: online-storage-and-backup, web-based-email, web-hosting, personal-sites-and-blogs, and content-delivery-networks.



In Security policy, block Quick UDP Internet Connections (QUIC) protocol *unless for business reasons, you want to allow encrypted browser traffic. Chrome and some other browsers establish sessions using QUIC instead of TLS/SSL, but QUIC uses proprietary encryption that the firewall can't decrypt, so potentially dangerous traffic may enter the network as encrypted traffic. Blocking QUIC forces the browser to fall back to TLS/SSL and enables the firewall to decrypt the traffic.*



Decrypting TLS traffic forces browsers that use HTTP/2 to fall back to HTTP 1.1 because the firewall can't decrypt HTTP/2 traffic. Allow browsers to fall back to HTTP 1.1 so you can decrypt this traffic and prevent potentially dangerous traffic from entering the network as encrypted traffic.

- Create a Decryption Profile
- Create a Decryption Policy Rule

Create a Decryption Profile

A decryption profile allows you to perform checks on both decrypted traffic and SSL traffic that you *choose* to exclude from decryption. (If a server breaks SSL decryption technically due to certificate pinning or other reasons, add the server to the Decryption Exclusion list.) Depending on your needs, create Decryption profiles to:

- Block sessions based on certificate status, including blocking sessions with expired certificates, untrusted issuers, unknown certificate status, certificate status check timeouts, and certificate extensions.
- Block sessions with unsupported versions and cipher suites, and that require using client authentication.
- Block sessions if the resources to perform decryption are not available or if a hardware security module is not available to sign certificates.
- Define the protocol versions and key exchange, encryption, and authentication algorithms allowed for SSL Forward Proxy and SSL Inbound Inspection traffic in the SSL Protocol Settings.

© 2018 Palo Alto Networks, Inc.

Don't weaken the main Decryption profile that you apply to most sites to accommodate weaker sites. Instead, create one or more separate Decryption profiles for sites that you need to support but that don't support strong ciphers and algorithms. You can also create different Decryption profiles for different URL Categories to fine tune security vs. performance for traffic that contains no sensitive material; however, you should always decrypt and inspect all the traffic you can.

After you create a decryption profile, attach it to a decryption policy rule; the firewall then enforces the decryption profile settings on traffic that matches the decryption policy rule.

Palo Alto Networks firewalls include a default decryption profile that you can use to enforce the basic recommended protocol versions and cipher suites for decrypted traffic. However, the best practice is to enable tighter decryption controls as described in SSL Forward Proxy Decryption Profile, SSL Inbound Inspection Decryption Profile, and SSL Protocol Settings Decryption Profile.



Avoid supporting weak protocols or algorithms because they contain known vulnerabilities that attackers can exploit. If you must allow a weaker protocol or algorithm to support a key partner or contractor who uses legacy systems with weak protocols, create a separate Decryption profile for the exception and attach it to a Decryption policy rule that applies the profile only to the relevant traffic (for example, the source IP address of the partner). Don't allow the weak protocol for all traffic.

STEP 1 | Create a new decryption profile.

Select **Objects** > **Decryption Profile**, **Add** or modify a decryption profile rule, and give the rule a descriptive **Name**.

STEP 2 | (Optional) Allow the profile rule to be **Shared** across every virtual system on a firewall or every Panorama device group.

STEP 3 | (Decryption Mirroring Only) Enable an Ethernet Interface for the firewall to use to copy and forward decrypted traffic.

Separate from this task, follow the steps to Configure Decryption Port Mirroring. Be aware of local privacy regulations that may prohibit mirroring or control the type of traffic that you can mirror. Decryption port mirroring requires a decryption port mirror license.

STEP 4 | (Optional) Block and control SSL tunneled and/or inbound traffic.



Although applying a Decryption profile to decrypted traffic is optional, it is a best practice to always apply a Decryption profile to the policy rules to protect your network against encrypted threats. You can't protect yourself against threats you can't see.

Select **SSL Decryption**:

- Select **SSL Forward Proxy** to configure the SSL Forward Proxy Decryption Profile settings to verify certificates, enforce protocol versions and cipher suites, and perform failure checks on SSL decrypted traffic. These settings are active only when this profile is attached to a decryption policy rule configured to perform SSL Forward Proxy decryption.
- Select **SSL Inbound Inspection** to configure the SSL Inbound Inspection Decryption Profile settings to enforce protocol versions and cipher suites and to perform failure checks on inbound SSL traffic. These settings are active only when this profile is attached to a decryption policy rule that performs SSL Inbound Inspection.
- Select **SSL Protocol Settings** to configure the SSL Protocol Settings Decryption Profile that control minimum and maximum protocol versions and key exchange, encryption, and authentication algorithms to enforce on decrypted SSL traffic. These settings are active when this profile is attached

© 2018 Palo Alto Networks, Inc.

to decryption policy rules that are set to perform either SSL Forward Proxy decryption or SSL Inbound Inspection.

STEP 5 | (Optional) Block and control traffic (for example, a URL category) for which you choose to Create a Policy-Based Decryption Exclusion.



Although applying a Decryption profile to traffic that you choose not to decrypt is optional, it is a best practice to always apply a Decryption profile to the policy rules to protect your network against sessions with expired certificates or untrusted issuers.

Select **No Decryption** to configure the Decryption Profile for No Decryption and check the **Block sessions with expired certificates** and **Block sessions with untrusted issuers** boxes to validate certificates for traffic that is excluded from decryption. Create policy-based exclusions only for traffic that you choose not to decrypt. If a server breaks decryption for technical reasons, don't create a policy-based exclusion, add the server to the SSL Decryption Exclusion list (**Device** > **Certificate Management** > **SSL Decryption Exclusion**).

These setting are active only when the decryption profile is attached to a decryption policy rule that disables decryption for certain traffic.

STEP 6 | (Optional) Block and control decrypted SSH traffic.

Select **SSH Proxy** to configure the SSH Proxy Decryption Profile and configure settings to enforce supported protocol versions and to block sessions if system resources are not available to perform decryption.

These settings are active only when the decryption profile is attached to a decryption policy rule that decrypts SSH traffic.

STEP 7 | Add the decryption profile when you Create a Decryption Policy Rule.

The firewall applies the decryption profile to and enforces the profile's settings on the traffic that matches the decryption policy rule.

STEP 8 | **Commit** the configuration.

Create a Decryption Policy Rule

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.

STEP 1 | Add a new decryption policy rule.

Select **Policies** > **Decryption**, **Add** a new decryption policy rule, and give the policy rule a descriptive **Name**.

STEP 2 | Configure the decryption rule to match to traffic based on network and Policy Objects:

- **Firewall security zones**—Select **Source** and/or **Destination** to match traffic based on the **Source Zone** and/or the **Destination Zone**.
- **IP addresses, address objects, and/or address groups**—Select **Source** and/or **Destination** to match traffic based on **Source Address** and/or the **Destination Address**. Alternatively, select **Negate** to exclude the source address list from decryption.
- **Users**—Select **Source** and set the **Source User** for whom to decrypt traffic. You can decrypt specific user or group traffic, or decrypt traffic for certain types of users, such as unknown users or pre-logon users (users that are connected to GlobalProtect but are not yet logged in).

© 2018 Palo Alto Networks, Inc.

- **Ports and protocols**—Select **Service/URL Category** to set the rule to match to traffic based on service. By default, the policy rule is set to decrypt **Any** traffic on TCP and UDP ports. You can **Add** a service or a service group, and optionally set the rule to **application-default** to match to applications only on the application default ports.



The application-default setting can be useful when you Create a Policy-Based Decryption Exclusion*. You can exclude applications running on their default ports from decryption, while continuing to decrypt the same applications when they are detected on non-standard ports.*

- **URLs and URL categories**—Select **Service/URL Category** and decrypt traffic based on:
  - An externally-hosted list of URLs that the firewall retrieves for policy-enforcement (see **Objects** > **External Dynamic Lists**).
  - Custom URL categories (see **Objects** > **Custom Objects** > **URL Category**). For example, you can create a custom URL Category to specify a group of sites you need to access for business purposes but that don't support the safest protocols and algorithms, and then apply a customized Decryption profile to allow the looser protocols and algorithms for just those sites (that way, you don't decrease security by downgrading the Decryption profile you use for most sites).
  - Palo Alto Networks predefined URL categories. This option is useful when you create policy-based decryption exclusions because you can exclude sensitive sites by category instead of individually. For example, although you can create a custom URL category to group sites that you do not want to decrypt, you can also exclude financial or healthcare-related sites from decryption based on the predefined Palo Alto Networks URL categories. You can also block risky URL Categories and create comfort pages to communicate the reason the sites are blocked or Enable Users to Opt Out of SSL Decryption.

STEP 3 | Set the rule to either decrypt matching traffic or to exclude matching traffic from decryption.

Select **Options** and set the policy rule **Action**:

To decrypt matching traffic:

1. Set the **Action** to **Decrypt**.
2. Set the **Type** of decryption for the firewall to perform on matching traffic:
   - SSL Forward Proxy
   - SSL Inbound Inspection. If you want to enable SSL Inbound Inspection, also select the **Certificate** for the destination internal server for the inbound SSL traffic.
   - SSH Proxy

To exclude matching traffic from decryption:

Set the **Action** to **No Decrypt**.

STEP 4 | (Optional) Select a **Decryption Profile** to perform additional checks on traffic that matches the policy rule.



Although applying a Decryption profile to decrypted traffic is optional, it is a best practice to always apply a Decryption profile to the policy rules to protect your network against encrypted threats. You can't protect yourself against threats you can't see.

For example, attach a decryption profile to a policy rule to ensure that server certificates are valid and to block sessions using unsupported protocols or ciphers. To Create a Decryption Profile, select **Objects** > **Decryption Profile**.

1. Create a decryption policy rule or open an existing rule to modify it.

© 2018 Palo Alto Networks, Inc.

2. Select **Options** and select a **Decryption Profile** to block and control various aspects of the traffic matched to the rule.

   The profile rule settings the firewall applies to matching traffic depends on the policy rule **Action** (Decrypt or No Decrypt) and the policy rule **Type** (SSL Forward Proxy, SSL Inbound Inspection, or SSH Proxy). This allows you to use the different decryption profiles with different types of decryption policy rules that apply to different types of traffic and users.
3. Click **OK**.

STEP 5 | Click **OK** to save the policy.

STEP 6 | Choose your next step to fully enable the firewall to decrypt traffic:

- Configure SSL Forward Proxy
- Configure SSL Inbound Inspection
- Configure SSH Proxy
- Create policy-based Decryption Exclusions for traffic you *choose* not to decrypt and add sites that break decryption for technical reasons such as pinned certificates or mutual authentication to the SSL Decryption Exclusion list.

© 2018 Palo Alto Networks, Inc.

Configure SSL Forward Proxy

To enable the firewall to perform 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. The firewall can use certificates signed by an enterprise certificate authority (CA) or self-signed certificates generated on the firewall as *Forward Trust certificates* to authenticate the SSL session with the client.

- (Recommended Best Practice) **Enterprise CA-signed Certificates**—An enterprise CA can issue a signing certificate that the firewall can use to sign the certificates for sites which require SSL decryption. When the firewall trusts the CA that signed the certificate of the destination server, the firewall can send a copy of the destination server certificate to the client, signed by the enterprise CA. This is a best practice because usually all network devices already trust the Enterprise CA (it is usually already installed in the devices' CA Trust storage), so you don't need to deploy the certificate on the endpoints, so the rollout process is smoother.
- **Self-signed Certificates**—The firewall can act as a CA and generate self-signed certificates that the firewall can use to sign the certificates for sites which require SSL decryption. The firewall can sign a copy of the server certificate to present to the client and establish the SSL session. This method requires that you need to install the self-signed certificates on all of your network devices so that those devices recognize the firewall's self-signed certificates. Because the certificates must be deployed to all devices, this method is better for small deployments and proof-of-concept (POC) trials than for large deployments.

Additionally, set up a *Forward Untrust certificate* for the firewall to present to clients when the server certificate is signed by a CA that the firewall does not trust. This ensures that clients are prompted with a certificate warning when attempting to access sites with untrusted certificates.



Regardless of whether you generate Forward Trust certificates from your Enterprise Root CA or use a self-signed certificate generated on the firewall, generate a separate subordinate Forward Trust CA certificate for each firewall. The flexibility of using separate subordinate CAs enables you to revoke one certificate when you decommission a device (or device pair) without affecting the rest of the deployment and reduces the impact in any situation in which you need to revoke a certificate. Separate Forward Trust CAs on each firewall also helps troubleshoot issues because the CA error message the user sees includes information about the firewall the traffic is traversing. If you use the same Forward Trust CA on every firewall, you lose the granularity of that information.

After setting up the Forward Trust and Forward Untrust certificates required for SSL Forward Proxy decryption, create a Decryption policy rule to define the traffic you want the firewall to decrypt and create a Decryption profile to apply SSL controls and checks to the traffic. The Decryption policy decrypts SSL tunneled traffic that matches the rule into clear text traffic. The firewall blocks and restricts traffic based on the Decryption profile attached to the Decryption policy and on the firewall Security policy. The firewall re-encrypts traffic as it exits the firewall.

STEP 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces.

View configured interfaces on the **Network** > **Interfaces** > **Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is.

© 2018 Palo Alto Networks, Inc.

STEP 2 | Configure the Forward Trust certificate for the firewall to present to clients when a trusted CA has signed the server certificate. You can use an enterprise CA-signed certificate or a self-signed certificate as the forward trust certificate.

(Recommended Best Practice) Use an enterprise CA-signed certificate as the Forward Trust certificate. Create a uniquely named Forward Trust certificate on each firewall:

1. Generate a Certificate Signing Request (CSR) for the enterprise CA to sign and validate:
    1. Select **Device** > **Certificate Management** > **Certificates** and click **Generate**.
    2. Enter a **Certificate Name**. Use a unique name for each firewall.
    3. In the **Signed By** drop-down, select **External Authority (CSR)**.
    4. (Optional) If your enterprise CA requires it, add **Certificate Attributes** to further identify the firewall details, such as Country or Department.
    5. Click **Generate** to save the CSR. The pending certificate is now displayed on the **Device Certificates** tab.
2. Export the CSR:
    1. Select the pending certificate displayed on the **Device Certificates** tab.
    2. Click **Export** to download and save the certificate file.

    Leave Export private key unselected in order to ensure that the private key remains securely on the firewall.
    3. Click **OK**.
3. Provide the certificate file to your enterprise CA. When you receive the enterprise CA-signed certificate from your enterprise CA, save the enterprise CA-signed certificate to import onto the firewall.
4. Import the enterprise CA-signed certificate onto the firewall:
    1. Select **Device** > **Certificate Management** > **Certificates** and click **Import**.
    2. Enter the pending **Certificate Name** exactly. The **Certificate Name** that you enter must exactly match the pending certificate name in order for the pending certificate to be validated.
    3. Select the signed **Certificate File** that you received from your enterprise CA.
    4. Click **OK**. The certificate is displayed as valid with the Key and CA check boxes selected.
5. Select the validated certificate to enable it as a **Forward Trust Certificate** to be used for SSL Forward Proxy decryption.
6. Click **OK** to save the enterprise CA-signed forward trust certificate.

Use a self-signed certificate as the Forward Trust certificate:

1. Create a self-signed Root CA certificate.
2. Click the self-signed root CA certificate (**Device** > **Certificate Management** > **Certificates** > **Device Certificates**) to open **Certificate information** and then click the **Trusted Root CA** checkbox.
3. Click **OK**.
4. Generate new subordinate CA certificates for each firewall:
    1. Select **Device** > **Certificate Management** > **Certificates**.
    2. Click **Generate** at the bottom of the window.
    3. Enter a **Certificate Name**.
    4. Enter a **Common Name**, such as 192.168.2.1. This should be the IP or FQDN that will appear in the certificate. In this case, we are using the IP of the trust interface. Avoid using spaces in this field.
    5. In the **Signed By** field, select the self-signed Root CA certificate that you created.
    6. Click the **Certificate Authority** check box to enable the firewall to issue the certificate. Selecting this check box creates a certificate authority (CA) on the firewall that is imported to the client browsers, so clients trust the firewall as a CA.

© 2018 Palo Alto Networks, Inc.

7. **Generate** the certificate.
5. Click the new certificate to modify it and click the **Forward Trust Certificate** checkbox to configure the certificate as the Forward Trust Certificate.
6. Click **OK** to save the self-signed forward trust certificate.
7. Repeat this procedure to generate a unique subordinate CA certificate on each firewall.

STEP 3 | Distribute the forward trust certificate to client system certificate stores.

If you are using an enterprise-CA signed certificate as the forward trust certificate for SSL Forward Proxy decryption, and the client systems already have the enterprise CA installed in the local trusted root CA list, you can skip this step. (The client systems trust the subordinate CA certificates you generate on the firewall because the Enterprise Trusted Root CA has signed them.)



If you do not install the forward trust certificate on client systems, users see certificate warnings for each SSL site they visit.

On a firewall configured as a GlobalProtect portal:



This option is supported with Windows and Mac client OS versions, and requires GlobalProtect agent 3.0.0 or later to be installed on the client systems.

1. Select **Network** > **GlobalProtect** > **Portals** and then select an existing portal configuration or **Add** a new one.
2. Select **Agent** and then select an existing agent configuration or **Add** a new one.
3. **Add** the self-signed firewall Trusted Root CA certificate to the Trusted Root CA section. After GlobalProtect distributes the firewall's Trusted Root CA certificate to client systems, the client systems trust the firewall's subordinate CA certificates because the clients trust the firewall's Root CA certificate.
4. **Install in Local Root Certificate Store** so that the GlobalProtect portal automatically distributes the certificate and installs it in the certificate store on GlobalProtect client systems.
5. Click **OK** twice.

Without GlobalProtect:

Export the firewall Trusted Root CA certificate so that you can import it into client systems. Highlight the certificate and click **Export** at the bottom of the window. Choose PEM format.



Do not select the Export private key checkbox! The private key should remain on the firewall and should not be exported to client systems.

Import the firewall's Trusted Root CA certificate into the browser Trusted Root CA list on the client systems in order for the clients to trust it. When importing into the client browser, ensure that you add the certificate to the Trusted Root Certification Authorities certificate store. On Windows systems, the default import location is the Personal certificate store. You can also simplify this process by using a centralized deployment option, such as an Active Directory Group Policy Object (GPO).

STEP 4 | Configure the Forward Untrust certificate (use the same Forward Untrust certificate for all firewalls).

1. Click **Generate** at the bottom of the certificates page.
2. Enter a **Certificate Name**, such as my-ssl-fwd-untrust.
3. Set the **Common Name**, for example 192.168.2.1. Leave **Signed By** blank.
4. Click the **Certificate Authority** check box to enable the firewall to issue the certificate.
5. Click **Generate** to generate the certificate.
6. Click **OK** to save.

© 2018 Palo Alto Networks, Inc.

7. Click the new my-ssl-fwd-untrust certificate to modify it and enable the **Forward Untrust Certificate** option.

   Do not export the Forward Untrust certificate to the Certificate Trust Lists of your network devices! Do not install the Forward Untrust certificate on client systems. This is critical because installing the Untrust certificate in the Trust List results in devices trusting websites that the firewall does not trust. In addition, users won't see certificate warnings for untrusted sites, so they won't know the sites are untrusted and may access those sites, which could expose your network to threats.
8. Click **OK** to save.

STEP 5 | (Optional) Configure the Key Size for SSL Forward Proxy Server Certificates that the firewall presents to clients. By default, the firewall determines the key size to use based on the key size of the destination server certificate.

STEP 6 | Create a Decryption Policy Rule to define traffic for the firewall to decrypt and Create a Decryption Profile to apply SSL controls to the traffic.

Although Decryption profiles are optional, it is a best practice to include a Decryption profile with each Decryption policy rule to prevent weak, vulnerable protocols and algorithms from allowing questionable traffic on your network.

1. Select **Policies** > **Decryption**, Add or modify an existing rule, and define traffic to be decrypted.
2. Select **Options** and:
   - Set the rule **Action** to **Decrypt** matching traffic.
   - Set the rule **Type** to **SSL Forward Proxy**.
   - (Optional but a best practice) Configure or select an existing **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, create a decryption profile to perform certificate checks and enforce strong cipher suites and protocol versions).
3. Click **OK** to save.

STEP 7 | Enable the firewall to forward decrypted SSL traffic for WildFire analysis.

This option requires an active WildFire license and is a WildFire best practice*.*

STEP 8 | **Commit** the configuration.

STEP 9 | Choose your next step:

- Enable Users to Opt Out of SSL Decryption.
- Configure Decryption Exclusions to disable decryption for certain types of traffic.

© 2018 Palo Alto Networks, Inc.

Configure SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL traffic destined for a network server (you can perform SSL Inbound Inspection for any server if you load the server certificate onto the firewall). With an SSL Inbound Inspection Decryption policy enabled, the firewall decrypts all SSL traffic identified by the policy to clear text traffic and inspects it. The firewall blocks, restricts, or allows the traffic based on the Decryption profile attached to the policy and the Security policy that applies to the traffic, including and any configured Antivirus, Vulnerability Protection, Anti-Spyware, URL-Filtering, and File Blocking profiles. As a best practice, enable the firewall to forward decrypted SSL traffic for WildFire analysis and signature generation.

Configuring SSL Inbound Inspection includes installing the targeted server certificate on the firewall, creating an SSL Inbound Inspection Decryption policy, and applying a Decryption profile to the policy.

STEP 1 | Ensure that the appropriate interfaces are configured as either Tap, Virtual Wire, Layer 2, or Layer 3 interfaces.



You cannot use a Tap mode interface for SSL Inbound Inspection if the negotiated cyphers include PFS key-exchange algorithms (DHE and ECDHE).

View configured interfaces on the **Network** > **Interfaces** > **Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including the interface type.

STEP 2 | Ensure that the targeted server certificate is installed on the firewall.

On the web interface, select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** to view device certificates installed on the firewall.

To import the targeted server certificate onto the firewall:

1. On the **Device Certificates** tab, select **Import**.
2. Enter a descriptive **Certificate Name**.
3. Browse for and select the targeted server **Certificate File**.
4. Click **OK**.

STEP 3 | Create a Decryption Policy Rule to define traffic for the firewall to decrypt and Create a Decryption Profile to apply SSL controls to the traffic.



Although Decryption profiles are optional, it is a best practice to include a Decryption profile with each Decryption policy rule to prevent weak, vulnerable protocols and algorithms from allowing questionable traffic on your network.

1. Select **Policies** > **Decryption**, **Add** or modify an existing rule, and define traffic to be decrypted.
2. Select **Options** and:
   - Set the rule **Action** to **Decrypt** matching traffic.
   - Set the rule **Type** to **SSL Inbound Inspection**.
   - Select the **Certificate** for the internal server that is the destination of the inbound SSL traffic.
   - (Optional but a best practice) Configure or select an existing **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, create a Decryption profile to terminate sessions with unsupported algorithms and unsupported cipher suites).

© 2018 Palo Alto Networks, Inc.

When you configure the SSL Protocol Settings Decryption Profile for SSL Inbound Inspection traffic, create separate profiles for servers with different security capabilities. For example, if one set of servers supports only RSA, the SSL Protocol Settings only need to support RSA. However, the SSL Protocol Settings for servers that support PFS should support PFS. Configure SSL Protocol Settings for the highest level of security that the server supports, but check performance to ensure that the firewall resources can handle the higher processing load that higher security protocols and algorithms require.

3. Click **OK** to save.

STEP 4 | Enable the firewall to forward decrypted SSL traffic for WildFire analysis.

This option requires an active WildFire license and is a WildFire best practice.

STEP 5 | **Commit** the configuration.

STEP 6 | Choose your next step:

- Enable Users to Opt Out of SSL Decryption.
- Configure Decryption Exclusions to disable decryption for certain types of traffic.

© 2018 Palo Alto Networks, Inc.

Configure SSH Proxy

Configuring SSH Proxy does not require certificates and the firewall generates the key used to decrypt SSH sessions automatically during boot up. With SSH decryption enabled, the firewall decrypts SSH traffic and blocks and or restricts the SSH traffic based on your decryption policy and decryption profile settings. Traffic is re-encrypted as it exits the firewall.

STEP 1 | Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. Decryption can only be performed on virtual wire, Layer 2, or Layer 3 interfaces.

View configured interfaces on the **Network** > **Interfaces** > **Ethernet** tab. The **Interface Type** column displays if an interface is configured to be a **Virtual Wire** or **Layer 2**, or **Layer 3** interface. You can select an interface to modify its configuration, including what type of interface it is.

STEP 2 | Create a Decryption Policy Rule to define traffic for the firewall to decrypt and Create a Decryption Profile to apply SSL controls to the traffic.



Although Decryption profiles are optional, it is a best practice to include a Decryption profile with each Decryption policy rule to prevent weak, vulnerable protocols and algorithms from allowing questionable traffic on your network.

1. Select **Policies** > **Decryption**, Add or modify an existing rule, and define traffic to be decrypted.
2. Select **Options** and:
   - Set the rule **Action** to **Decrypt** matching traffic.
   - Set the rule **Type** to **SSH Proxy**.
   - (Optional but a best practice) Configure or select an existing **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, create a Decryption profile to terminate sessions with unsupported versions and unsupported algorithms).
3. Click **OK** to save.

STEP 3 | **Commit** the configuration.

STEP 4 | Choose your next step:

- Enable Users to Opt Out of SSL Decryption.
- Configure Decryption Exclusions to disable decryption for certain types of traffic.

© 2018 Palo Alto Networks, Inc.

Configure Server Certificate Verification for Undecrypted Traffic

You create no-decryption policies for traffic that you *choose* not to decrypt because the traffic is personal, sensitive, or subject to local laws and regulations. For example, you may choose not to decrypt the traffic of certain executives or traffic between finance users and finance servers that contain personal information. (Don't exclude traffic that you can't decrypt because a site breaks decryption for technical reasons such as a pinned certificate or mutual authentication by policy. Instead, add the hostname to the Decryption Exclusion List.)

However, just because you don't decrypt the traffic doesn't mean you should let any and all undecrypted traffic on your network. It is a best practice to apply a No Decryption profile to undecrypted traffic to block sessions with expired certificates and untrusted issuers.

STEP 1 | Create a Decryption Policy Rule to identify the undecrypted traffic and Create a Decryption Profile to block bad sessions.

1. Select **Policies** > **Decryption** and Add or modify an existing rule to identify the undecrypted traffic.
2. Select **Options** and:
   - Set the rule **Action** to **No Decrypt** so that the firewall doesn't decrypt traffic that matches the rule.
   - Ignore the rule **Type** because the traffic is not decrypted.
   - (Optional but a best practice) Configure or select an existing Decryption Profile for No Decryption to block sessions with expired certificates and untrusted certificate issuers.

STEP 2 | **Commit** the configuration.

STEP 3 | Choose your next step:

- Enable Users to Opt Out of SSL Decryption.
- Configure Decryption Exclusions to disable decryption for certain types of traffic.

© 2018 Palo Alto Networks, Inc.

Decryption Exclusions

You can exclude two types of traffic from decryption:

- Traffic that breaks decryption for *technical reasons*, such as using a pinned certificate or mutual authentication. Palo Alto Networks provides a predefined SSL Decryption Exclusion list (**Device** > **Certificate management** > **SSL Decryption Exclusion**) that excludes hosts with applications and services that are known to break decryption technically from SSL Decryption by default. If you encounter sites that break decryption technically and are not on the SSL Decryption Exclusion list, you can add them to list manually by server hostname. The firewall blocks sites whose applications and services break decryption technically unless you add them to the SSL Decryption Exclusion list.
- Traffic that you *choose* not to decrypt because of business, regulatory, personal, or other reasons, such as financial, health, military, or government traffic. You can choose to exclude traffic based on source, destination, URL category, and service.

To increase visibility into traffic and reduce the attack surface as much as possible, don't make decryption exceptions unless you must.

- Palo Alto Networks Predefined Decryption Exclusions
- Exclude a Server from Decryption for Technical Reasons
- Create a Policy-Based Decryption Exclusion

Palo Alto Networks Predefined Decryption Exclusions

The firewall provides a predefined SSL Decryption Exclusion list to exclude from decryption commonly used sites that 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 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 Security policy on it.

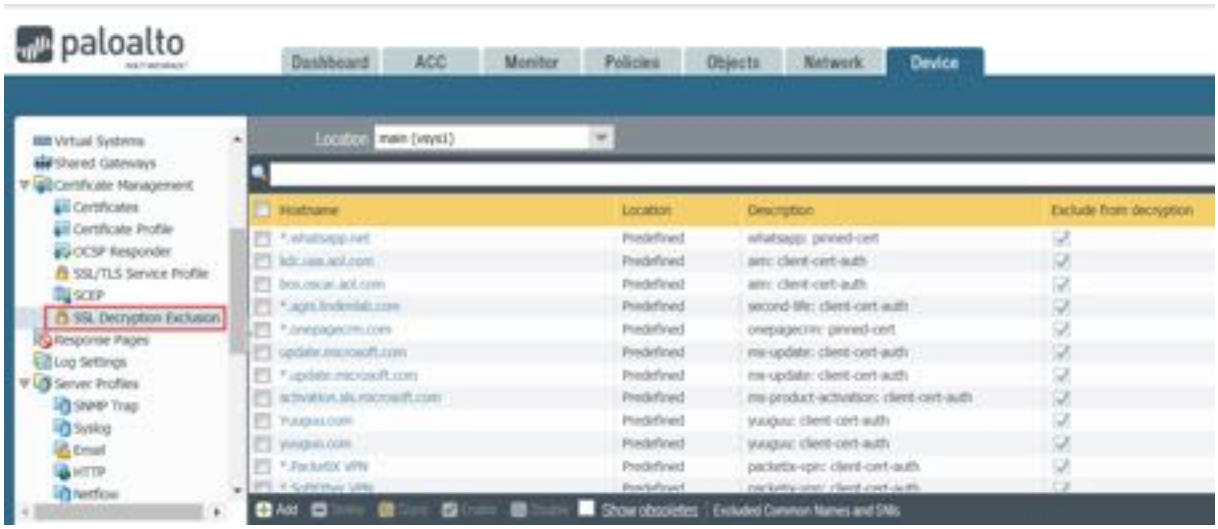


*The SSL Decryption Exclusion list is **not** for sites that you choose not to decrypt for legal, regulatory, business, privacy, or other volitional reasons, it is only for sites that break decryption technically. For traffic (IP addresses, users, URL categories, services, and even entire zones) that you choose not to decrypt,* Create a Policy-Based Decryption Exclusion*.*

Because the traffic of sites on the SSL Decryption Exclusion list remains encrypted, the firewall does not inspect or provide further security enforcement the traffic. You can disable a predefined exclusion. For example, you may choose to disable predefined exclusions to enforce a strict security policy that allows only applications and services that the firewall can inspect and on which the firewall can enforce Security policy. However, the firewall blocks 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 Palo Alto Networks predefined SSL decryption exclusions directly on the firewall (**Device** > **Certificate Management** > **SSL Decryption Exclusions**).

© 2018 Palo Alto Networks, Inc.



The **Hostname** displays the name of the host that houses the application or service that breaks decryption technically. You can also **Add** hosts to Exclude a Server from Decryption for Technical Reasons if it is not on the predfined list.

The **Description** displays the reason the firewall can't decrypt the site's traffic, for example, **pinned-cert** (a pinned certificate) or **client-cert-auth** (client authentication).

The firewall automatically removes enabled predefined SSL decryption exclusions from the list when they become obsolete (the firewall removes an application that decryption previously caused to break when the application becomes supported with decryption). **Show Obsoletes** checks if any disabled predefined exclusions remain on the list and are no longer needed. The firewall does not remove disabled predefined decryption exclusions from the list automatically, but you can select and **Delete** obsolete entries.

You can select a hostname's checkbox and then click **Disable** to remove predefined sites from the list. Use the SSL Decryption Exclusion list only for sites that break decryption for technical reasons, don't use it for sites that you choose not to decrypt.

Exclude a Server from Decryption for Technical Reasons

If decryption breaks an important application or service technically, you can add the hostname of the site that hosts to the application or service to the Palo Alto Networks predefined SSL Decryption Exclusion list to create a custom decryption exception. The firewall doesn't decrypt, inspect, and enforce Security policy on traffic that the SSL Decryption Exclusion list allows because the traffic remains encrypted, so be sure that the sites you add to the list really are sites with applications or services you need for business. For example, some business-critical internal custom applications may break decryption and you can add them to the list so that the firewall allows the encrypted custom application traffic.

*The SSL Decryption Exclusion list is **not** for sites that you choose not to decrypt for legal, regulatory, business, privacy, or other volitional reasons, it is only for sites that break decryption technically. For traffic (IP addresses, users, URL categories, services, and even entire zones) that you choose not to decrypt,* Create a Policy-Based Decryption Exclusion*.*

Reasons that sites break decryption technically include pinned certificates, mutual authentication, incomplete certificate chains, and unsupported ciphers.

If the technical reason for excluding a site from decryption is an incomplete certificate chain, the next-generation firewall doesn't automatically fix the chain as a browser would. If you need to add a site to the SSL Decryption Exclusion list, manually review the site to ensure

© 2018 Palo Alto Networks, Inc.

it's a legitimate business site, then download the missing sub-CA certificates and load and deploy *them onto the firewall.*

After you add a server to the SSL Decryption Exclusion list, the firewall compares the server hostname that you use to define the decryption exclusion against the common name (CN) in the certificate a server presents. If a single server hosts multiple websites using different certificates, the firewall compares the hostname against the server name indication (SNI) that the client presents to indicate the server to which it wants to connect.

STEP 1 | Select **Device** > **Certificate Management** > **SSL Decryption Exclusions**.

STEP 2 | **Add** a new decryption exclusion, or select an existing custom entry to modify it.

STEP 3 | Enter the **hostname** of the website or application you want to exclude from decryption.

To exclude all hostnames associated with a certain domain from decryption, you can use a wildcard asterisk (*). In this case, all sessions where the server presents a CN that contains the domain are excluded from decryption.

Make sure that the hostname field is unique for each custom entry. If a predefined decryption exclusion hostname matches the hostname of a custom entry, the custom entry takes precedence.

STEP 4 | (Optional) Select **Shared** to share the exclusion across all virtual systems in a multiple virtual system firewall.

STEP 5 | **Exclude** the application from decryption.

The following example shows an SSL Decryption exclusion for all hostnames associated with the fictitious site *pinnedcerts.com*, shares the exception with all virtual systems on the firewall, provides the reason for the exclusion (the **Description**, which displays in the **Description** column of the SSL Decryption Exclusion list), and excludes the traffic from decryption.



Alternatively, if you modify an existing decryption exclusion, you can clear the **Exclude** checkbox to start decrypting traffic from that host. However, if the traffic still breaks decryption technically, the firewall blocks it.

STEP 6 | Click **OK** to save the new SSL Decryption Exclusion list entry.

Create a Policy-Based Decryption Exclusion

Policy-based decryption exclusions are for excluding traffic that you *choose* not to decrypt. You can create a policy-based decryption exclusion based on any combination of the traffic's source, destination, service, or URL Category. Examples of traffic you may choose not to decrypt include:

- Sensitive traffic that you should never decrypt, such as the URL Filtering categories financial-services, health-and-medicine, government, and military.

© 2018 Palo Alto Networks, Inc.

You can use predefined URL Categories to except entire categories from decryption, you can create custom URL Categories to define a customized list of URLs that you don't want to decrypt, or you can create an External Dynamic List (EDL) to define a customized list of URLs that you don't want to decrypt.



In environments such as Office 365 that have dynamically changing IP addresses or in environments where you make frequent changes to the list of URLs that you want to exclude from decryption, it's often preferable to use an EDL instead of a URL Category to specify the excluded URLs. Editing a custom URL Category or adding or removing a predefined URL Category requires a Commit action to take effect. However, when you edit an EDL, the firewall imports and implements the changes dynamically, without a Commit action, so using an EDL is less disruptive in a dynamically changing environment.

- The traffic originates or is destined for executives or other users whose traffic shouldn't be decrypted.
- Some devices such as finance servers may need to be excepted from decryption.
- Depending on the business, some companies may value privacy and the user experience more than security for some applications.
- Laws or local regulations that prohibit decryption of some traffic.

  An example of not decrypting traffic for regulatory and legal compliance is the European Union (EU) General Data Protection Regulation (GDPR). The EU GDPR will require strong protection of all personal data for all individuals. The GDPR affects all companies, including foreign companies, that collect or process the personal data of EU residents.

  Different regulations and compliance rules may mean that you treat the same data differently in different countries or regions. Businesses usually can decrypt personal information in their corporate data centers because the business owns the information. The best practice is to decrypt as much traffic as possible so that you can see it and apply security protection to it.

If you choose to except traffic from decryption, make sure it really is traffic you don't want to decrypt. Don't create exceptions that are broader than necessary. Configure the traffic source, destination, service, and URL Category in the Decryption policy so that you except only the traffic you really mean to except from decryption. The more specific the decryption exclusion, the better, so that you don't inadvertently exclude more traffic than necessary from decryption. Work with human resources, legal, finance, executives, IT, and other groups to identify sensitive information that the firewall should not decrypt.



Similar to Security policy rules, the firewall compares incoming traffic to Decryption policy rules in the Decryption policy rulebase's sequence. Place exclusion policy rules at the top of the rulebase.

The reason you place decryption exclusion rules at the top of the rulebase is to prevent inadvertently decrypting traffic you don't want to decrypt or traffic that laws or regulations prevent you from decrypting because the traffic matches a decryption rule before it matches the rule that is meant to except the traffic from decryption. If you create policy-based decryption exclusions, the best practice is to place the following exclusion rules at the top of the decryption rulebase, in the following order:

1. IP-address based exceptions for sensitive destination servers.
2. Source-user based exceptions for executives and other users or groups.
3. Custom URL or EDL based exceptions for destination URLs.
4. Sensitive predefined URL Category based exceptions for destination URLs of entire categories such as financial services, health and medicine, government, military, shopping, etc.

Place rules that decrypt traffic after these rules in the decryption rulebase.

© 2018 Palo Alto Networks, Inc.

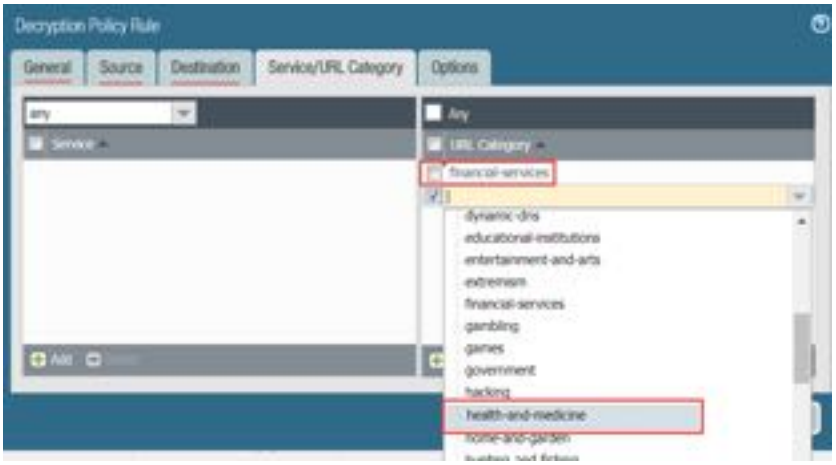
STEP 1 | Exclude traffic from decryption based on match criteria.

This example shows how to exclude traffic categorized as financial or health-related from SSL Forward Proxy decryption.

1. Select **Policies** > **Decryption** and **Add** or modify a decryption policy rule.
2. Define the traffic that you want to exclude from decryption.

   In this example:

   1. Give the rule a descriptive **Name**, such as No-Decrypt-Finance-Health.
   2. Set the **Source** and **Destination** to **Any** to apply the No-Decrypt-Finance-Health rule to all SSL traffic destined for an external server.
   3. Select **URL Category** and **Add** the URL categories financial-services and health-and-medicine.



3. Select **Options** and set the rule to **No Decrypt**.
4. (Optional but a best practice) Use a Decryption Profile for No Decryption to validate certificates for sessions the firewall does not decrypt. Attach the No Decryption profile to the rule and set the profile to **Block sessions with expired certificates** and **Block sessions with untrusted issuers**.
5. Click **OK** to save the No-Decrypt-Finance-Health decryption rule.

STEP 2 | Place the decryption exclusion rule at the top of your decryption policy rulebase.

The firewall enforces decryption rules against incoming traffic in the rulebase sequence and enforces the first rule that match the traffic.

Select the **No-Decrypt-Finance-Health** policy (**Decryption** > **Policies**), and click **Move Up** until it appears at the top of the list, or drag and drop the rule.

STEP 3 | Save the configuration.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Enable Users to Opt Out of SSL Decryption

In privacy-sensitive situations, you may want to alert your users that the firewall is decrypting certain web traffic and allow them either to continue to the site with the understanding that their traffic is decrypted or to terminate the session and be block from going to the site. (There is no option to go to the site and also avoid decryption.)

The first time a user attempts to browse to an HTTPS site or application that matches the decryption policy, the firewall displays a response page notifying users that it will decrypt the session. Users can either click **Yes** to allow decryption and continue to the site or click **No** to opt out of decryption and terminate the session. The choice to allow decryption applies to all HTTPS sites that users try to access for the next 24 hours, after which the firewall redisplays the response page. Users who opt out of SSL decryption cannot access the requested web page, or any other HTTPS site, for the next minute. After the minute elapses, the firewall redisplays the response page the next time the users attempt to access an HTTPS site.

The firewall includes a predefined SSL Decryption Opt-out Page that you can enable. You can optionally customize the page with your own text and/or images. However, the best practice is not to allow users to opt out of decryption if you don't have to.



Custom response pages larger than the maximum supported size are not decrypted or displayed to users. In PAN-OS 8.0.11 and earlier releases, custom response pages on a decrypted site cannot exceed 8,191 bytes; the maximum size is increased to 17,999 bytes in PAN-OS 8.0.12 and later PAN-OS 8.0 releases.

STEP 1 | (Optional) Customize the SSL Decryption Opt-out Page.

1. Select **Device** > **Response Pages**.
2. Select the **SSL Decryption Opt-out Page** link.
3. Select the **Predefined** page and click **Export**.
4. Using the HTML text editor of your choice, edit the page.
5. If you want to add an image, host the image on a web server that is accessible from your end user systems.
6. Add a line to the HTML to point to the image. For example:

```
<img src="http://cdn.slidesharecdn.com/ Acme-logo-96x96.jpg?1382722588"/>
```

7. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding.
8. Back on the firewall, select **Device** > **Response Pages**.
9. Select the **SSL Decryption Opt-out Page** link.
10. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.
11. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select shared to make it available to all virtual systems.
12. Click **OK** to import the file.
13. Select the response page you just imported and click **Close**.

STEP 2 | Enable SSL Decryption Opt Out.

1. On the **Device** > **Response Pages** page, click the **Disabled** link.
2. Select the **Enable SSL Opt-out Page** and click **OK**.
3. **Commit** the changes.

STEP 3 | Verify that the Opt Out page displays when you attempt to browse to a site.

© 2018 Palo Alto Networks, Inc.

From a browser, go to an encrypted site that matches your decryption policy.

Verify that the SSL Decryption Opt-out response page displays.



© 2018 Palo Alto Networks, Inc.

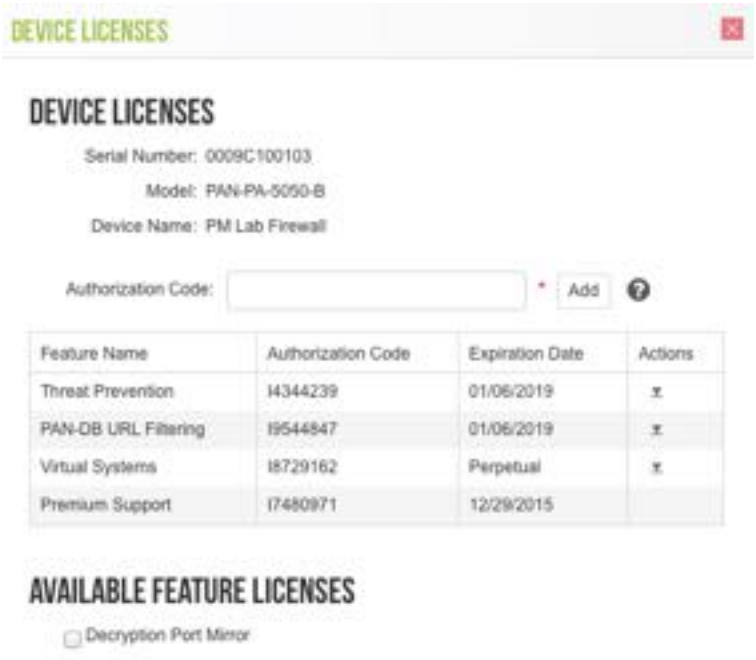
Configure Decryption Port Mirroring

Before you can enable Decryption Mirroring, you must obtain and install a Decryption Port Mirror license. The license is free of charge and can be activated through the support portal as described in the following procedure. After you install the Decryption Port Mirror license and reboot the firewall, you can enable decryption port mirroring.

Keep in mind that the decryption, storage, inspection, and/or use of SSL traffic is regulated in certain countries and user consent may be required in order to use the decryption mirror feature. Additionally, use of this feature could enable malicious users with administrative access to the firewall to harvest usernames, passwords, social security numbers, credit card numbers, or other sensitive information submitted using an encrypted channel. Palo Alto Networks recommends that you consult with your corporate counsel before activating and using this feature in a production environment.

STEP 1 | Request a license for each firewall on which you want to enable decryption port mirroring.

1. Log in to the Palo Alto Networks Customer Support website and navigate to the **Assets** tab.
2. Select the entry for the firewall you want to license and select **Actions**.
3. Select **Decryption Port Mirror**. A legal notice displays.
4. If you are clear about the potential legal implications and requirements and still want to set up decryption port mirroring, click **I understand and wish to proceed**.
5. Click **Activate**.



STEP 2 | Install the Decryption Port Mirror license on the firewall.

1. From the firewall web interface, select **Device** > **Licenses**.
2. Click **Retrieve license keys from license server**.
3. Verify that the license has been activated on the firewall.

© 2018 Palo Alto Networks, Inc.



4. Reboot the firewall (**Device** > **Setup** > **Operations**). This feature is not available for configuration until PAN-OS reloads.

STEP 3 | Enable the firewall to forward decrypted traffic. Superuser permission is required to perform this step.

On a firewall with a single virtual system:

1. Select **Device** > **Setup** > **Content - ID**.
2. Select the **Allow forwarding of decrypted content** check box.
3. Click **OK** to save.

On a firewall with multiple virtual systems:

1. Select **Device** > **Virtual System.**
2. Select a Virtual System to edit or create a new Virtual System by selecting **Add**.
3. Select the **Allow forwarding of decrypted content** check box.
4. Click **OK** to save.

STEP 4 | Enable an Ethernet interface to be used for decryption mirroring.

1. Select **Network** > **Interfaces** > **Ethernet**.
2. Select the Ethernet interface that you want to configure for decryption port mirroring.
3. Select **Decrypt Mirror** as the **Interface Type**.

   This interface type will appear only if the Decryption Port Mirror license is installed.
4. Click **OK** to save.

STEP 5 | Enable mirroring of decrypted traffic.

1. Select **Objects** > **Decryption Profile**.
2. Select an **Interface** to be used for **Decryption Mirroring**.

   The **Interface** drop-down contains all Ethernet interfaces that have been defined as the type: **Decrypt Mirror**.
3. Specify whether to mirror decrypted traffic before or after policy enforcement.

   By default, the firewall mirrors all decrypted traffic to the interface before security policy lookup, which allows you to replay events and analyze traffic that generates a threat or triggers a drop action. If you want to only mirror decrypted traffic after security policy enforcement, select the **Forwarded Only** check box. With this option, only traffic that is forwarded through the firewall is mirrored. This option is useful if you are forwarding the decrypted traffic to other threat detection devices, such as a DLP device or another intrusion prevention system (IPS).
4. Click **OK** to save the decryption profile.

STEP 6 | Attach the decryption profile (with decryption port mirroring enabled) to a decryption policy rule. All traffic decrypted based on the policy rule is mirrored.

1. Select **Policies** > **Decryption**.
2. Click **Add** to configure a decryption policy or select an existing decryption policy to edit.
3. In the **Options** tab, select **Decrypt** and the **Decryption Profile** created in step 4.
4. Click **OK** to save the policy.

© 2018 Palo Alto Networks, Inc.

STEP 7 | Save the configuration.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Verify Decryption

After you configure a best practice Decryption profile and apply it to traffic in a Decryption policy rule, check the log files to verify that the firewall is decrypting the traffic that you intend to decrypt and that the firewall is not decrypting the traffic that you don't want to decrypt. In addition, follow post-deployment decryption best practices to maintain the deployment.

- **View Decrypted Traffic Sessions**—Filter the Traffic Logs (**Monitor** > **Logs** > **Traffic**) using the filter **`( flags has proxy )`**.

  This filter displays only logs in which the SSL proxy flag is on, meaning only decrypted traffic—every log entry has the value yes in the **Decrypted** column.



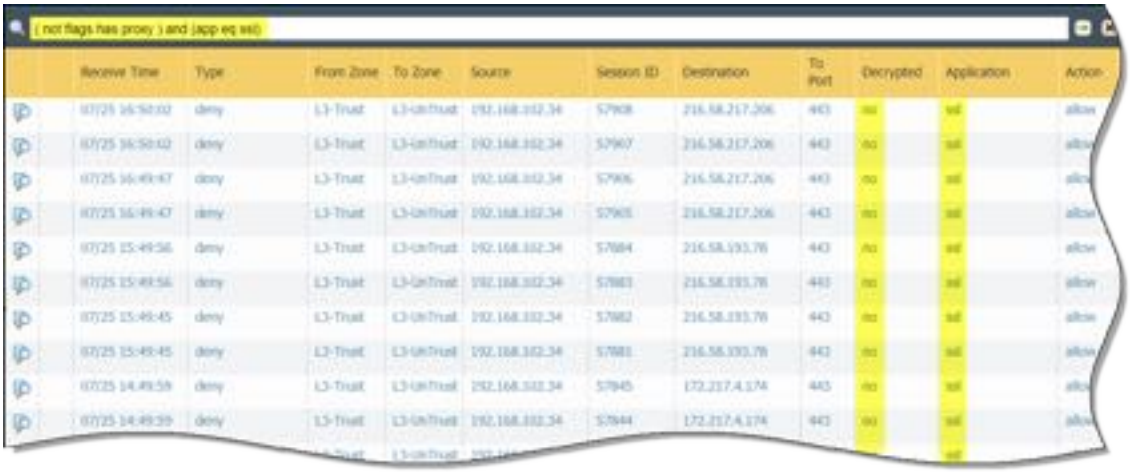
  You can filter the traffic in a more granular fashion by adding more terms to the filter. For example, you can filter for decrypted traffic going only to the destination IP address 172.217.3.206 by adding the filter **`( addr.dst in 172.217.3.206 )`**:



- **View SSL Traffic Sessions That Are Not Decrypted**—Filter the Traffic Logs (**Monitor** > **Logs** > **Traffic**) using the filter **`( not flags has proxy ) and ( app eq ssl )`**.

  This filter displays only logs in which the SSL proxy flag is off (meaning only encrypted traffic) and the traffic is SSL traffic; every log entry has the value **no** in the **Decrypted** column and the value **ssl** in the **Application** column.

© 2018 Palo Alto Networks, Inc.



Similar to the example for viewing decrypted traffic logs, you can add terms to filter the traffic that you don't decrypt in a more granular fashion.

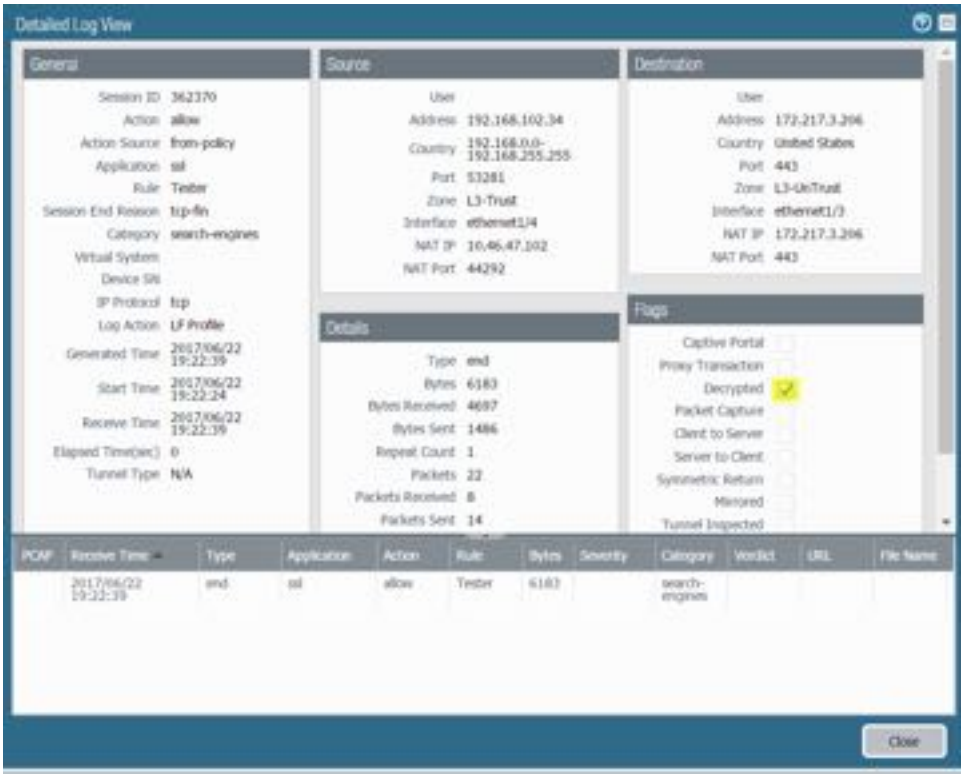
- **View The Log for a Particular Session**—To view the decryption log for a particular session, filter on the Session ID.

  For example, to see the log for a session with the ID 362370, filter using the term **`( sessionid eq 362370 )`**. You can find the ID number in the Session ID column in the log output, as shown in the previous screens. If the Session ID column isn't displayed, add the column to the output.



- **Drill Down Into the Details**—To view more information about a particular log entry, click the magnifying glass to see a detailed log view. For example, for Session ID 362370 (shown in the previous bullet), the detailed log looks like this:

© 2018 Palo Alto Networks, Inc.



The box for the **Decrypted** flag provides a second way to verify if traffic was decrypted.

You can also take upstream and downstream packet captures of decrypted traffic to view how the firewall processes SSL traffic and takes actions on packets, or perform deep packet inspection.

© 2018 Palo Alto Networks, Inc.

Temporarily Disable SSL Decryption

In some cases you may want to temporarily disable SSL decryption. For example, if you deployed SSL decryption too hastily and something doesn't work correctly but you're not sure what it is, and you have a lot of rules to examine, you can use the CLI to temporarily turn off decryption and give yourself time to analyze and solve the issue. After solving the issue, you can use the CLI to turn SSL decryption back on again. Because temporarily disabling and then re-enabling decryption using the CLI doesn't require a Commit operation, you can do it without disrupting network traffic.

The following CLI commands temporarily disable SSL decryption without a Commit and re-enable decryption without a Commit.



The command to disable SSL decryption doesn't persist in the configuration after a reboot. If you turn off decryption temporarily and then reboot the firewall, regardless of whether the issue has been fixed, decryption is turned on again.

- Disable SSL Decryption

```
set system setting
ssl-decrypt skip-ssl-decrypt yes
```

- Re-enable SSL Decryption

```
set system setting
ssl-decrypt skip-ssl-decrypt no
```


© 2018 Palo Alto Networks, Inc.

URL Filtering

The Palo Alto Networks URL filtering solution allows you to monitor and control the sites users can access, to prevent phishing attacks by controlling the sites to which users can submit valid corporate credentials, and to enforce safe search for search engines like Google and Bing.

> URL Filtering Overview
> URL Filtering Concepts
> PAN-DB Categorization
> Enable a URL Filtering Vendor
> Determine URL Filtering Policy Requirements
> Configure URL Filtering
> Use an External Dynamic List in a URL Filtering Profile
> Customize the URL Filtering Response Pages
> Allow Password Access to Certain Sites
> Safe Search Enforcement
> Monitor Web Activity
> Set Up the PAN-DB Private Cloud
> URL Filtering Use Cases
> Troubleshoot URL Filtering

© 2018 Palo Alto Networks, Inc.

URL Filtering Overview

The Palo Alto Networks URL filtering solution compliments App-ID by enabling you to configure the firewall to identify and control access to web (HTTP and HTTPS) traffic and to protect your network from attack.

With URL Filtering enabled, all web traffic is compared against the URL filtering database, which contains a listing of millions of websites that have been categorized into categories. You can use these URL categories as a match criteria to enforce security policy and to 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 to Prevent Credential Phishing based on URL category.

Although the Palo Alto Networks URL filtering solution supports both BrightCloud and PAN-DB, only the PAN-DB URL filtering solution allows you to choose between the *PAN-DB Public Cloud* and the *PAN-DB Private Cloud*. Use the public cloud solution if the Palo Alto Networks next-generation firewalls on your network can directly access the Internet. If the network security requirements in your enterprise prohibit the firewalls from directly accessing the Internet, you can deploy a PAN-DB private cloud on one or more M-500 appliances that function as PAN-DB servers within your network.

- URL Filtering Vendors
- Interaction Between App-ID and URL Categories
- PAN-DB Private Cloud

URL Filtering Vendors

Palo Alto Networks firewalls support two URL filtering vendors:

- **PAN-DB**—A Palo Alto Networks developed URL filtering database that is tightly integrated into PAN-OS and the Palo Alto Networks threat intelligence cloud. PAN-DB provides high-performance local caching for maximum inline performance on URL lookups, and offers coverage against malicious URLs and IP addresses. As WildFire, which is a part of the Palo Alto Networks threat intelligence cloud, identifies unknown malware, zero-day exploits, and advanced persistent threats (APTs), the PAN-DB database is updated with information on malicious URLs so that you can block malware downloads, and disable Command and Control (C2) communications to protect your network from cyber threats. The URL categories malware and phishing are updated every five minutes, to ensure that you can manage access to these sites within minutes of categorization.

  To view a list of PAN-DB URL filtering categories, refer to https://urlfiltering.paloaltonetworks.com/CategoryList.aspx.
- **BrightCloud**—A third-party URL database that is owned by Webroot, Inc. that is integrated into PAN-OS firewalls. For information on the BrightCloud URL database, visit http://brightcloud.com.

For instructions on configuring the firewall to use one of the supported URL Filtering vendors, see Enable a URL Filtering Vendor.

Interaction Between App-ID and URL Categories

The Palo Alto Networks URL filtering solution in combination with App-ID on page 509 provides unprecedented protection against a full spectrum of cyber attacks, legal, regulatory, productivity, and resource utilization risks. While App-ID gives you control over what applications users can access, URL filtering provides control over related web activity. When combined with User-ID, you can enforce controls based on users and groups.

With today's application landscape and the way many applications use HTTP and HTTPS, you will need to use App-ID, URL filtering, or both in order to define comprehensive web access policies. App-ID signatures are granular and they allow you to identify shifts from one web-based application to another; URL filtering allows you to enforce actions based on a specific website or URL category. For example, while you can use

© 2018 Palo Alto Networks, Inc.

URL filtering to control access to Facebook and/or LinkedIn, URL filtering cannot block the use of related applications such as email, chat, or other any new applications that are introduced after you implement policy. When combined with App-ID, you can control the use of related applications because of the granular application signatures that can identify each application and regulate access to Facebook while blocking access to Facebook chat, when defined in policy.

You can also use URL categories as a match criteria in policies. Instead of creating policies limited to either allow all or block all behavior, URL as a match criteria permits exception-based behavior and gives you more granular policy enforcement capabilities. For example, deny access to malware and hacking sites for all users, but allow access to users that belong to the IT-security group.

For some examples, see URL Filtering Use Cases on page 691.

PAN-DB Private Cloud

The PAN-DB private cloud is an on-premise solution that is suitable for organizations that prohibit or restrict the use of the PAN-DB public cloud service. With this on-premise solution, you can deploy one or more M-500 appliances as PAN-DB servers within your network or data center. The firewalls query the PAN-DB private cloud to perform URL lookups, instead of accessing the PAN-DB public cloud.

The process for performing URL lookups, in both the private and the public cloud is the same for the firewalls on the network. By default, the firewall is configured to access the public PAN-DB cloud. If you deploy a PAN-DB private cloud, you must configure the firewalls with a list of IP addresses or FQDNs to access the server(s) in the private cloud.



Firewalls running PAN-OS 5.0 or later versions can communicate with the PAN-DB private cloud.

When you Set Up the PAN-DB Private Cloud on page 686, you can either configure the M-500 appliance(s) to have direct internet access or keep it completely offline. Because the M-500 appliance requires database and content updates to perform URL lookups, if the appliance does not have an active internet connection, you must manually download the updates to a server on your network and then, import the updates using SCP into each M-500 appliance in the PAN-DB private cloud. In addition, the appliances must be able to obtain the seed database and any other regular or critical content updates for the firewalls that it services.

To authenticate the firewalls that connect to the PAN-DB private cloud, a set of default server certificates are packaged with the appliance; you cannot import or use another server certificate for authenticating the firewalls. If you change the hostname on the M-500 appliance, the appliance automatically generates a new set of certificates to authenticate the firewalls.

- M-500 Appliance for PAN-DB Private Cloud on page 644
- Table 3: Differences Between the PAN-DB Public Cloud and PAN-DB Private Cloud on page 645

M-500 Appliance for PAN-DB Private Cloud

To deploy a PAN-DB private cloud, you need one or more M-500 appliances. The M-500 appliance ships in Panorama mode, and to be deployed as PAN-DB private cloud you must set it up to operate in PAN-URL-DB mode. In the PAN-URL-DB mode, the appliance provides URL categorization services for enterprises that do not want to use the PAN-DB public cloud.

The M-500 appliance when deployed as a PAN-DB private cloud uses two ports- MGT (Eth0) and Eth1; Eth2 is not available for use. The management port is used for administrative access to the appliance and for obtaining the latest content updates from the PAN-DB public cloud or from a server on your network. For communication between the PAN-DB private cloud and the firewalls on the network, you can use the MGT port or Eth1.

© 2018 Palo Alto Networks, Inc.

The M-100 appliance cannot be deployed as a PAN-DB private cloud.

The M-500 appliance in PAN-URL-DB mode:

- Does not have a web interface, it only supports a command-line interface (CLI).
- Cannot be managed by Panorama.
- Cannot be deployed in a high availability pair.
- Does not require a URL Filtering license. The firewalls, must have a valid PAN-DB URL Filtering license to connect with and query the PAN-DB private cloud.
- Ships with a set of default server certificates that are used to authenticate the firewalls that connect to the PAN-DB private cloud. You cannot import or use another server certificate for authenticating the firewalls. If you change the hostname on the M-500 appliance, the appliance automatically generates a new set of certificates to authenticate the firewalls that it services.
- Can be reset to Panorama mode only. If you want to deploy the appliance as a dedicated Log Collector, switch to Panorama mode and then set it in log collector mode.

Table 3: Differences Between the PAN-DB Public Cloud and PAN-DB Private Cloud

Differences	PAN-DB Public Cloud	PAN-DB Private Cloud
Content and Database Updates	Content (regular and critical) updates and full database updates are published multiple times during the day. The PAN-DB public cloud updates the URL categories malware and phishing every five minutes. The firewall checks for critical updates whenever it queries the cloud servers for URL lookups.	Content updates and full URL database updates are available once a day during the work week.
URL Categorization Requests	Submit URL categorization change requests using the following options: • Palo Alto Networks Test A Site website. • URL filtering profile setup page on the firewall. • URL filtering log on the firewall.	Submit URL categorization change requests only using the Palo Alto Networks Test A Site website.
Unresolved URL Queries	If the firewall cannot resolve a URL query, the request is sent to the servers in the public cloud.	If the firewall cannot resolve a query, the request is sent to the M-500 appliance(s) in the PAN-DB private cloud. If there is no match for the URL, the PAN-DB private cloud sends a category *unknown* response to the firewall; the request is not sent to the public cloud unless you have configured the M-500 appliance to access the PAN-DB public cloud. If the M-500 appliance(s) that constitute your PAN-DB private cloud is configured to be completely offline, it does not send any data or analytics to the public cloud.

© 2018 Palo Alto Networks, Inc.

URL Filtering Concepts

- URL Categories on page 646
- URL Filtering Profile on page 647
- URL Filtering Profile Actions on page 647
- URL Category Exception Lists on page 649
- External Dynamic List for URLs on page 651
- Container Pages on page 651
- HTTP Header Logging on page 652
- URL Filtering Response Pages on page 652
- URL Category as Policy Match Criteria on page 655

URL Categories

Each website defined in the URL filtering database is assigned a URL category. Here are a few ways to leverage URL categories:

- **Block or allow traffic based on URL category**—You can create a URL Filtering profile that specifies an action for each URL category and attach the profile to a policy. Traffic that matches the policy would then be subject to the URL filtering settings in the profile. For example, to block all gaming websites you would set the block action for the URL category *games* in the URL profile and attach it to the security policy rule(s) that allow web access. See Configure URL Filtering for more information.
- **Enforce policy based on URL category**—If you want a specific policy rule to apply only to web traffic to sites in a specific category, use the site URL category as match criteria when you create the policy rule. For example, you could use the URL category *streaming-media* in a QoS policy to apply bandwidth controls to all websites that are categorized as streaming media. See URL Category as Policy Match Criteria for more information.
- **Block or allow corporate credential submissions based on URL category**—Prevent Credential Phishing by enabling the firewall to detect corporate credential submissions to sites, and then block or allow those submissions based on URL category. Block users from submitting credentials to malicious and untrusted sites, warn users against entering corporate credentials on unknown sites or warn them against reusing corporate credentials on non-corporate sites, and explicitly allow users submit credentials to corporate and sanctioned sites.

By grouping websites into categories, it makes it easy to define actions based on certain types of websites. In addition to the standard URL categories, there are three additional categories:

Category	Description
not-resolved	Indicates that the website was not found in the local URL filtering database and the firewall was unable to connect to the cloud database to check the category. When a URL category lookup is performed, the firewall first checks the dataplane cache for the URL; if no match is found, it checks the management plane cache, and if no match is found there, it queries the URL database in the cloud. In the case of the PAN-DB private cloud, the URL database in the cloud is not used for queries. Setting the action to block for traffic that is categorized as not-resolved, may be very disruptive to users. You could set the action as continue, so that users you can notify users that they are accessing a site that is blocked by company policy and provide the option to read the disclaimer and continue to the website.

© 2018 Palo Alto Networks, Inc.

Category	Description
	For more information on troubleshooting lookup issues, see Troubleshoot URL Filtering.
private-ip-addresses	Indicates that the website is a single domain (no sub-domains), the IP address is in the private IP range, or the URL root domain is unknown to the cloud.
unknown	The website has not yet been categorized, so it does not exist in the URL filtering database on the firewall or in the URL cloud database. When deciding on what action to take for traffic categorized as *unknown*, be aware that setting the action to block may be very disruptive to users because there could be a lot of valid sites that are not in the URL database yet. If you do want a very strict policy, you could block this category, so websites that do not exist in the URL database cannot be accessed. Palo Alto Networks collects the list of URLs from the unknown category and processes them to determine the URL category. These URLs are processed automatically, everyday, provided the websites has machine readable content that is in a supported format and language. Upon categorization, the updated category information is made available to all PAN-DB customers. See Configure URL Filtering.



You can submit URL categorization change requests using the Palo Alto Networks dedicated web portal (Test A Site), the URL filtering profile setup page on the firewall, or the URL filtering log on the firewall. Each change request is automatically processed everyday, provided the websites provides machine readable content that is in a supported format and language. Sometimes, the categorization change requires a member of the Palo Alto Networks engineering staff to perform a manual review. In such cases, the process may take a little longer.

URL Filtering Profile

A URL filtering profile is a collection of URL filtering controls that you can apply 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. 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. You can then customize the newly-added URL profiles and add lists of specific websites that should always be blocked or allowed. For example, you may want to block social-networking sites, but allow some websites that are part of the social-networking category.



Configure a best practice URL Filtering profile to ensure protection against URLs that have been observed hosting malware or exploitive content.

URL Filtering Profile Actions

The URL Filtering profile specifies web access and credential submission permissions for each URL category. By default, site access for all URL categories is set to allow when you Create a new URL Filtering profile. This means that the users will be able to browse to all sites freely and the traffic will not be logged. You can customize the URL Filtering profile with custom **Site Access** settings for each category, or use the predefined default URL filtering profile on the firewall to allow access to all URL categories except the following threat-prone categories, which it blocks: abused-drugs, adult, gambling, hacking, malware, phishing, questionable, and weapons.

© 2018 Palo Alto Networks, Inc.

For each URL category, select the **User Credential Submissions** to allow or disallow users from submitting valid corporate credentials to a URL in that category in order to Prevent Credential Phishing. Managing the sites to which users can submit credentials requires User-ID and you must first Set Up Credential Phishing Prevention. URL categories with the **Site Access** set to block are automatically set to also block user credential submissions.



Learn more about configuring a best practice URL Filtering profile to ensure protection against URLs that have been observed hosting malware or exploitive content.

Action	Description
Site Access	
alert	The website is allowed and a log entry is generated in the URL filtering log.
allow	The website is allowed and no log entry is generated.
block	The website is blocked and the user will see a response page and will not be able to continue to the website. A log entry is generated in the URL filtering log. Blocking site access for a URL category also sets User Credential Submissions for that URL category to block.
continue	The user will be prompted with a response page indicating that the site has been blocked due to company policy, but the user is prompted with the option to continue to the website. The **continue** action is typically used for categories that are considered benign and is used to improve the user experience by giving them the option to continue if they feel the site is incorrectly categorized. The response page message can be customized to contain details specific to your company. A log entry is generated in the URL filtering log.  *The Continue page doesn't display properly on client systems configured to use a proxy server.*
override	The user will see a response page indicating that a password is required to allow access to websites in the given category. With this option, the security admin or helpdesk person would provide a password granting temporary access to all websites in the given category. A log entry is generated in the URL filtering log. See Allow Password Access to Certain Sites.  *The Override page doesn't display properly on client systems configured to use a proxy server.*
none	The **none** action only applies to custom URL categories. Select **none** to ensure that if multiple URL profiles exist, the custom category will not have any impact on other profiles. For example, if you have two URL profiles and the custom URL category is set to **block** in one profile, if you do not want the block action to apply to the other profile, you must set the action to **none**. Also, in order to delete a custom URL category, it must be set to **none** in any profile where it is used.

© 2018 Palo Alto Networks, Inc.

Action	Description
User Credential Permissions *These settings require you to first* Set Up Credential Phishing Prevention.	
alert	Allow users to submit corporate credentials to sites in this URL category, but generate a URL Filtering alert log each time this occurs.
allow (default)	Allow users to submit corporate credentials to websites in this URL category.
block	Block users from submitting corporate credentials to websites in this category. A default anti-phishing response page is displayed to users when they access sites to which corporate credential submissions are blocked. You can choose to create a custom block page to display.
continue	Display a response page to users that prompts them to select Continue to access to access the site. By default, the Anti Phishing Continue Page is shown to user when they access sites to which credential submissions are discouraged. You can also choose to create a custom response page to display—for example, if you want to warn users against phishing attempts or reusing their credentials on other websites.

URL Category Exception Lists

You can exclude specific websites from URL category enforcement, ensuring that these websites are blocked or allowed regardless of their associated URL category. For example, you could block a URL category but choose to allow certain sites that fall within that category. To create these kind of exceptions to URL category enforcement:

- Add the IP addresses or URLs of the sites you want to block or allow (regardless of their associated URL category) directly to a URL Filtering profile (**Objects** > **Security Profiles** > **URL Filtering** > **Overrides**.
- Use an External Dynamic List in a URL Filtering profile. The benefit to using an External Dynamic List to specify the sites you want to enforce separately from their URL categories is that you can update the External Dynamic List without performing a configuration change or commit on the firewall.

The following guidelines describe how to populate URL Category block and allow lists, or a text file that you're using as the source of an External Dynamic List for URLs:

- Basic Guidelines For URL Category Exception Lists
- Wildcard Guidelines for URL Category Exception Lists
- URL Category Exception List—Wildcard Examples

Basic Guidelines For URL Category Exception Lists

- Enter the IP addresses or URLs of websites that you want to enforce separately from the associated URL category.
- List entries must be an exact match and are case-insensitive.
- You can enter a string that is an exact match to the website (and possibly, specific subdomain) for which you want to control access, or you can use wildcard character to allow an entry to match to more than one website subdomains. For details on using wildcard characters, review Wildcard Guidelines for URL Category Exception Lists.

© 2018 Palo Alto Networks, Inc.

- Omit `http` and `https` from URL entries.

Wildcard Guidelines for URL Category Exception Lists

You can use wildcards in URL Category exception lists to easily configure a single entry to match to multiple website subdomains and pages, without having to specify exact subdomains and pages.

Follow these guidelines when creating wildcard entries:

- The following characters are considered token separators: . / ? & = ; +

  Every string separated by one or two of these characters is a token. Use wildcard characters as token placeholders, indicating that a specific token can contain any value.
- You can use either an asterisk (*) or a caret (^) in place of a token, to indicate a wildcard value.
- Wildcard characters must be the only character within a token; however, an entry can contain multiple wildcards.

When to use asterisk (*) wildcards:

Use an asterisk (*) wildcard to indicate one or multiple variable subdomains. For example, to specify enforcement for Palo Alto Network's website regardless of the domain extension used, which might be one or two subdomains depending on location, you would add the entry: `www.paloaltonetworks.*`. This entry would match to both www.paloaltonetworks.com and www.paloaltonetworks.co.uk.

When to use caret (^) wildcards:

Use caret (^) wildcards to indicate one variable subdomain, and might be helpful when targeting an exact number of subdomains for enforcement. For example, `mail.^.com`matches only to URLs like `mail.company.com`. This entry wouldn't match to a site like `mail.company.sso.com`, where the URL includes an additional subdomain.

Do not create an entry with consecutive asterisk (*) wildcards or more than nine consecutive caret (^) wildcards—entries like these can affect firewall performance.

For example, do not add an entry like `mail.*.*.com`; instead, depending on the range of websites you want to control access to, enter `mail.*.com` or `mail.^.^.com`. An entry like `mail.*.com` matches to a greater number of sites than`mail.^.^.com`; `mail.*.com` matches to sites with any number of subdomains and `mail.^.^.com` matches to sites with exactly two subdomains.

URL Category Exception List—Wildcard Examples

The following tables lists examples of URL exception lists entries using wildcards, and examples of the sites that these entries match to.

URL Exception List Entry	Matching Sites
Example Set 1	
*.company.com	eng.tools.company.com support.tools.company.com tools.company.com docs.company.com
^.company.com	tools.company.com docs.company.com

© 2018 Palo Alto Networks, Inc.

URL Exception List Entry	Matching Sites
^.^.company.com	eng.tools.company.com support.tools.company.com
Example Set 2	
mail.google.*	mail.google.com mail.google.co.uk
mail.google.^	mail.google.com
mail.google.^.^	mail.google.co.uk

Do not create an entry with consecutive asterisk (*) wildcards or more than nine consecutive caret (^) wildcards—entries like these can affect firewall performance.

For example, do not add an entry like `mail.*.*.com`; instead, depending on the range of websites you want to control access to, enter `mail.*.com` or `mail.^.^.com`. An entry like `mail.*.com` matches to a greater number of sites than `mail.^.^.com`; `mail.*.com` matches to sites with any number of subdomains and `mail.^.^.com` matches to sites with exactly two subdomains.

External Dynamic List for URLs

To protect your network from new sources of threat or malware, you can use External Dynamic List on page 1011 in URL Filtering profiles to block or allow, or to define granular actions such as continue, alert, or override for URLs, before you attach the profile to a Security policy rule. Unlike the allow list, block list, or a custom URL category on the firewall, an external dynamic list gives you the ability to update the list without a configuration change or commit on the firewall. The firewall dynamically imports the list at the configured interval and enforces policy for the URLs (IP addresses or domains will be ignored) in the list. For URL formatting guidelines, see URL Category Exception Lists on page 649.

Container Pages

A container page is the main page that a user accesses when visiting a website, but additional websites may be loaded within the main page. If the **Log Container page only** option is enabled in the URL filtering profile, only the main container page will be logged, not subsequent pages that may be loaded within the container page. Because URL filtering can potentially generate a lot of log entries, you may want to turn on this option, so log entries will only contain those URIs where the requested page file name matches the specific mime-types. The default set includes the following mime-types:

- application/pdf
- application/soap+xml
- application/xhtml+xml
- text/html
- text/plain
- text/xml

If you have enabled the Log container page only option, there may not always be a correlated URL log entry for threats detected by antivirus or vulnerability protection.

© 2018 Palo Alto Networks, Inc.

HTTP Header Logging

URL filtering provides visibility and control over web traffic on your network. For improved visibility into web content, you can configure the URL Filtering profile to log HTTP header attributes included in a web request. When a client requests a web page, the HTTP header includes the user agent, referer, and x-forwarded-for fields as attribute-value pairs and forwards them to the web server. When enabled for logging HTTP headers, the firewall logs the following attribute-value pairs in the URL Filtering logs:

Attribute	Description
User-Agent	The web browser that the user used to access the URL, for example, Internet Explorer. This information is sent in the HTTP request to the server.
Referer	The URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested.
X-Forwarded-For (XFF)	The option in the HTTP request header field that preserves the IP address of the user who requested the web page. If you have a proxy server on your network, the XFF allows you to identify the IP address of the user who requested the content, instead of only recording the proxy server's IP address as source IP address that requested the web page.

URL Filtering Response Pages

The firewall provides three predefined response pages that display by default when a user attempts to browse to a site in a category that is configured with one of the block actions in the URL Filtering Profile (block, continue, or override) or when Container Pages is enabled:

- **URL Filtering and Category Match Block Page**

  Access blocked by a URL Filtering Profile or because the URL category is blocked by a Security policy rule.



- **URL Filtering Continue and Override Page**

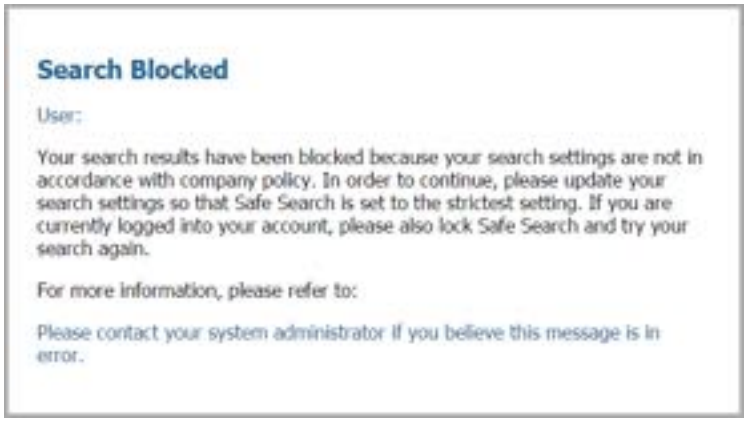
  Page with initial block policy that allows users to bypass the block by clicking **Continue**. With URL Admin Override enabled, (Allow Password Access to Certain Sites), after clicking **Continue**, the user must supply a password to override the policy that blocks the URL.

© 2018 Palo Alto Networks, Inc.



- **URL Filtering Safe Search Block Page**

  Access blocked by a Security policy rule with a URL Filtering profile that has the Safe Search Enforcement option enabled (see Safe Search Enforcement). The user will see this page if a search is performed using Google, Bing, Yahoo, or Yandex and their browser or search engine account setting for Safe Search is not set to strict.



- **Anti Phishing Block Page**

  This page displays to users when they attempt to enter corporate credentials (usernames or passwords) on a web page in a category for which credential submissions are blocked. The user can continue to access the site but remains unable to submit valid corporate credentials to any associated web forms. To control the sites to which users can submit corporate credentials, the firewall must be configured with User-ID and enabled to To Prevent Credential Phishing based on URL category.



- **Anti Phishing Continue Page**

  This page warns users against submitting credentials (usernames and passwords) to a web site. Warning users against submitting credentials can help to discourage them from reusing corporate credentials and to educate them about possible phishing attempts. They must select Continue to proceed to credentials on the site. To control the sites to which users can submit corporate credentials, the firewall must be configured with User-ID and enabled to To Prevent Credential Phishing based on URL category.

© 2018 Palo Alto Networks, Inc.



You can either use the predefined pages, or you can Customize the URL Filtering Response Pages to communicate your specific acceptable use policies and/or corporate branding. In addition, you can use the URL Filtering Response Page Variables for substitution at the time of the block event or add one of the supported Response Page References to external images, sounds, or style sheets.

Table 4: URL Filtering Response Page Variables

Variable	Usage
<user/>	The firewall replaces the variable with the username (if available via User-ID) or IP address of the user when displaying the response page.
<url/>	The firewall replaces the variable with the requested URL when displaying the response page.
<category/>	The firewall replaces the variable with the URL filtering category of the blocked request.
<pan_form/>	HTML code for displaying the **Continue** button on the URL Filtering Continue and Override page.

You can also add code that triggers the firewall to display different messages depending on what URL category the user is attempting to access. For example, the following code snippet from a response page specifies to display Message 1 if the URL category is games, Message 2 if the category is travel, or Message 3 if the category is kids:

```
var cat = "<category/>";
switch(cat)
{
 case 'games':
 document.getElementById("warningText").innerHTML = "Message 1";
 break;
 case 'travel':
 document.getElementById("warningText").innerHTML = "Message 2";
 break;
 case 'kids':
 document.getElementById("warningText").innerHTML = "Message 3";
 break;
}
```


© 2018 Palo Alto Networks, Inc.

Only a single HTML page can be loaded into each virtual system for each type of block page. However, other resources such as images, sounds, and cascading style sheets (CSS files) can be loaded from other servers at the time the response page is displayed in the browser. All references must include a fully qualified URL.

Table 5: Response Page References

Reference Type	Example HTML Code
Image	`<img src="http://virginiadot.org/images/Stop-Sign-gif.gif">`
Sound	`<embed src="http://simplythebest.net/sounds/WAV/WAV_files/ movie_WAV_files/ do_not_go.wav" volume="100" hidden="true" autostart="true">`
Style Sheet	`<link href="http://example.com/style.css" rel="stylesheet" type="text/css" />`
Hyperlink	`<a href="http://en.wikipedia.org/wiki/Acceptable_use_policy">View Corporate Policy</a>`

URL Category as Policy Match Criteria

Use URL Categories as a match criteria in a policy rule for more granular enforcement. For example, suppose you have configured Decryption, but you want to exclude traffic to certain types of websites (for example, healthcare or financial services) from being decrypted. In this case you could create a decryption policy rule that matches those categories and set the action to no-decrypt. By placing this rule above the rule to decrypt all traffic, you can ensure that web traffic with URL categories that match the no-decrypt rule, and all other traffic would match the subsequent rule.

The following table describes the policy types that accept URL category as match criteria:

Policy Type	Description
Authentication	To ensure that users authenticate before being allowed access to a specific category, you can attach a URL category as a match criterion for Authentication policy rules.
Decryption	Decryption policies can use URL categories as match criteria to determine if specified websites should be decrypted or not. For example, if you have a decryption policy with the action decrypt for all traffic between two zones, there may be specific website categories, such as *financial-services* and/or *health-and-medicine*, that should not be decrypted. In this case, you would create a new decryption policy with the action of *no#decrypt* that precedes the decrypt policy and then defines a list of URL categories as match criteria

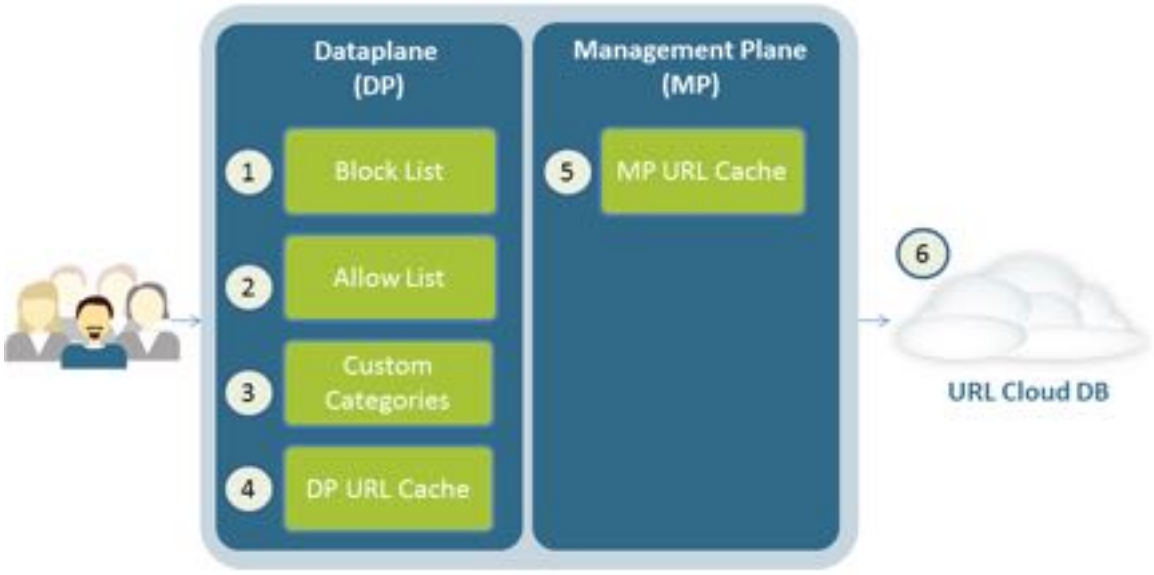
© 2018 Palo Alto Networks, Inc.

Policy Type	Description
	for the policy. By doing this, each URL category that is part of the no-decrypt policy will not be decrypted. You could also configure a custom URL category to define your own list of URLs that can then be used in the no-decrypt policy.
QoS	QoS policies can use URL categories to allocate throughput levels for specific website categories. For example, you may want to allow the streaming-media category, but limit throughput by adding the URL category as match criteria to the QoS policy.
Security	In security policies you can use URL categories both as a match criteria in the **Service/URL Category** tab, and in URL filtering profiles that are attached in the **Actions** tab. If for example, the IT-security group in your company needs access to the hacking category, while all other users are denied access to the category, you must create the following rules: • A Security policy rule that allows the IT-Security group to access content categorized as hacking. The Security policy rule references the *hacking* category in the S**ervices/URL Category** tab and IT-Security group in the **Users** tab. • Another Security policy rule that allows general web access for all users. To this rule you attach a URL filtering profile that blocks the hacking category. The policy that allows access to hacking must be listed before the policy that blocks hacking. This is because security policy rules are evaluated top down, so when a user who is part of the security group attempts to access a hacking site, the policy rule that allows access is evaluated first and will allow the user access to the hacking sites. Users from all other groups are evaluated against the general web access rule which blocks access to the hacking sites.

© 2018 Palo Alto Networks, Inc.

PAN-DB Categorization

When a user requests a URL the firewall determines the URL category by comparing the URL with the following components (in order) until it finds a match:



URL Categorization Resolution Process Flow

If a requested URL matches an expired entry in the dataplane (DP) URL cache, the cache responds with the expired category, but also sends a URL categorization query to the management plane (MP) cache. This prevents unnecessary delays in the DP, assuming that the frequency of category change is low. Similarly, in the MP URL cache, if a URL query from the DP cache matches an expired entry in the MP cache, the MP responds to the DP with the expired category and will also send a URL categorization request to the PAN-DB cloud database. Upon getting the response from the cloud, the firewall sends the updated category to the DP.

As new URLs and categories are defined or if critical updates are needed, the cloud database is updated. Each time the firewall queries the cloud for a URL lookup or if no cloud lookups have occurred for 30 minutes, the database versions on the firewall be compared and if they do not match, an incremental update will be performed.

The following table describes the PAN-DB components in detail. The BrightCloud system works similarly, but does not use an initial seed database.

Component	Description
URL Filtering Seed Database	The initial seed database downloaded to the firewall is a small subset of the database that is maintained on the Palo Alto Networks URL cloud servers. The reason this is done is because the full database contains millions of URLs and many of these URLs may never be accessed by your users. When downloading the initial seed database, you select a region (North America, Europe, APAC, Japan). Each region contains a subset of URLs most accessed for the given region. This allows the firewall to store a much smaller URL database for better URL lookup performance. If a user accesses a website that is not in the local URL database, the firewall queries the full cloud database and then adds the new URL to the local database. This way the local database

© 2018 Palo Alto Networks, Inc.

Component	Description
	on the firewall is continually populated/customized based on actual user activity. *Re-downloading the PAN-DB seed database or switching the URL database vendor from PAN-DB to BrightCloud will clear the local database.*
Cloud Service See Differences Between the PAN-DB Public Cloud and PAN-DB Private Cloud, for information on the private cloud.	The PAN-DB cloud service is implemented using Amazon Web Services (AWS). AWS provides a distributed, high-performance, and stable environment for seed database downloads and URL lookups for Palo Alto Networks firewalls and communication is performed over SSL. The AWS cloud systems hold the entire PAN-DB and is updated as new URLs are identified. The PAN-DB cloud service supports an automated mechanism to update the local URL database on the firewall if the version does not match. Each time the firewall queries the cloud servers for URL lookups, it will also check for critical updates. If there have been no queries to the cloud servers for more than 30 minutes, the firewall will check for updates on the cloud systems. The cloud system also provides a mechanism to submit URL category change requests. This is performed through the test-a-site service and is available directly from the firewall (URL filtering profile setup) and from the Palo Alto Networks Test A Site website. You can also submit a URL categorization change request directly from the URL filtering log on the firewall in the log details section.
Management Plane (MP) URL Cache	When you activate PAN-DB on the firewall, the firewall downloads a seed database from one of the PAN-DB cloud servers to initially populate the local cache for improved lookup performance. Each regional seed database contains the top URLs for the region and the size of the seed database (number of URL entries) also depends on the platform. The URL MP cache is automatically written to the local drive on the firewall every eight hours, before the firewall is rebooted, or when the cloud upgrades the URL database version on the firewall. After rebooting the firewall, the file that was saved to the local drive will be loaded to the MP cache. A least recently used (LRU) mechanism is also implemented in the URL MP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs.
Dataplane (DP) URL Cache	This is a subset of the MP cache and is a customized, dynamic URL database that is stored in the dataplane (DP) and is used to improve URL lookup performance. The URL DP cache is cleared at each firewall reboot. The number of URLs that are stored in the URL DP cache varies by hardware platform and the current URLs stored in the TRIE (data structure). A least recently used (LRU) mechanism is implemented in the DP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs. Entries in the URL DP cache expire after a specified period of time; this expiration period is not configurable.

© 2018 Palo Alto Networks, Inc.

Enable a URL Filtering Vendor

To enable URL filtering on a firewall, you must purchase and activate a URL Filtering license for one of the supported URL Filtering Vendors and then install the database for the vendor you selected.

Starting with PAN-OS 6.0, firewalls managed by Panorama do not need to be running the same URL filtering vendor that is configured on Panorama. For firewalls running PAN-OS 6.0 or later, when a mismatch is detected between the vendor enabled on the firewalls and what is enabled on Panorama, the firewalls can automatically migrate URL categories and/or URL profiles to (one or more) categories that align with that of the vendor enabled on it. For guidance on how to configure URL Filtering on Panorama if you are managing firewalls running different PAN-OS versions, refer to the Panorama Administrator's Guide.

If you have valid licenses for both PAN-DB and BrightCloud, activating the PAN-DB license automatically deactivates the BrightCloud license (and vice versa). At a time, only one URL filtering license can be active on a firewall.

- Enable PAN-DB URL Filtering
- Enable BrightCloud URL Filtering

Enable PAN-DB URL Filtering

STEP 1 | Obtain and install a PAN-DB URL filtering license and confirm that it is installed.

If the license expires, PAN-DB URL Filtering continues to work based on the URL category information that exists in the dataplane and management plane caches. However, URL cloud lookups and other cloud#based updates will not function until you install a valid license.

1. Select **Device** > **Licenses** and, in the License Management section, select the license installation method:
   - **Retrieve license keys from license server**
   - **Activate feature using authorization code**
   - **Manually upload license key**
2. After installing the license, confirm that the PAN-DB URL Filtering section, **Date Expires** field, displays a valid date.



STEP 2 | Download the initial seed database and activate PAN-DB URL Filtering.

The firewall must have Internet access; you cannot manually upload the PAN-DB seed database.

1. In the PAN-DB URL Filtering section, **Download Status** field, click **Download Now**.
2. Choose a region and then click **OK** to start the download.
3. After the download completes, click **Activate**. The value in the Active field changes to Yes.

© 2018 Palo Alto Networks, Inc.



If PAN-DB is already the active URL filtering vendor, clicking Re-Download clears the dataplane and management plane caches and replaces them with a new seed database. You should avoid doing this unless it is necessary, as you will lose your cache, which is customized based on your users' web traffic.

STEP 3 | Schedule the firewall to download dynamic updates for Applications and Threats.

A Threat Prevention license is required to receive content updates, which covers Antivirus and Applications and Threats.

1. Select **Device** > **Dynamic Updates**.
2. In the Schedule field in the Applications and Threats section, click the **None** link to schedule periodic updates.

   You can only schedule dynamic updates if the firewall has direct internet access. If updates are already scheduled in a section, the link text displays the schedule settings.

   The Applications and Threats updates sometimes contain updates for URL filtering related to Safe Search Enforcement.

Enable BrightCloud URL Filtering

STEP 1 | Obtain and install a BrightCloud URL filtering license and confirm that it is installed.

BrightCloud has an option in the URL filtering profile (**Objects** > **Security Profiles** > **URL Filtering**) to either allow all categories or block all categories if the license expires.

1. Select **Device** > **Licenses** and, in the **License Management** section, select the license installation method:
   - **Activate feature using authorization code**
   - **Retrieve license keys from license server**
   - **Manually upload license key**
2. After installing the license, confirm that the BrightCloud URL Filtering section, **Date Expires** field, displays a valid date.

STEP 2 | Install the BrightCloud database.

The way you do this depends on whether or not the firewall has direct Internet access.

Firewall with Direct Internet Access

Select **Device** > **Licenses** and in the BrightCloud URL Filtering section, Active field, click the **Activate** link to install the BrightCloud database. This operation automatically initiates a system reset.

Firewall without Direct Internet Access

1. Download the BrightCloud database to a host that has Internet access. The firewall must have access to the host:

© 2018 Palo Alto Networks, Inc.

1. On a host with Internet access, go to the Palo Alto Networks Customer Support web site, www.paloaltonetworks.com/support/tabs/overview.html, and log in.
2. In the Resources section, click **Dynamic Updates**.
3. In the BrightCloud Database section, click **Download** and save the file to the host.

2. Upload the database to the firewall:
   1. Log in to the firewall, select **Device** > **Dynamic Updates** and click **Upload**.
   2. For the **Type**, select **URL Filtering**.
   3. Enter the path to the **File** on the host or click **Browse** to find it, then click **OK**. When the Status is **Completed**, click **Close**.
3. Install the database:
   1. Select **Device** > **Dynamic Updates** and click **Install From File**.
   2. For the **Type**, select **URL Filtering**. The firewall automatically selects the file you just uploaded.
   3. Click **OK** and, when the Result is **Succeeded**, click **Close**.

STEP 3 | Enable cloud lookups for dynamically categorizing a URL if the category is not available on the local BrightCloud database.

1. Access the PAN-OS CLI.
2. Enter the following commands to enable dynamic URL filtering:

```
> configure
# set deviceconfig setting url dynamic-url yes
# commit
```

STEP 4 | Schedule the firewall to download dynamic updates for Applications and Threats signatures and URL filtering.

You can only schedule dynamic updates if the firewall has direct Internet access.

The Applications and Threats updates might contain updates for URL filtering related to the **Safe Search Enforcement** option in the URL filtering profile. For example, if Palo Alto Networks adds support for a new search provider vendor or if the method used to detect the Safe Search setting for an existing vendor changes, the Application and Threats updates will include that update.

BrightCloud updates include a database of approximately 20 million websites that are stored locally on the firewall. You must schedule URL filtering updates to receive BrightCloud database updates.



A Threat Prevention license is required to receive Antivirus and Applications and Threats updates.

1. Select **Device** > **Dynamic Updates**.
2. In the Applications and Threats section, Schedule field, click the **None** link to schedule periodic updates.
3. In the URL Filtering section, Schedule field, click the **None** link to schedule periodic updates.



If updates are already scheduled in a section, the link text displays the schedule settings.

© 2018 Palo Alto Networks, Inc.

Determine URL Filtering Policy Requirements

The recommended practice for deploying URL filtering in your organization is to first start with a passive URL filtering profile that will alert on most categories. After setting the alert action, you can then monitor user web activity for a few days to determine patterns in web traffic. After doing so, you can then make decisions on the websites and website categories that should be controlled.

In the procedure that follows, threat#prone sites will be set to block and the other categories will be set to alert, which will cause all websites traffic to be logged. This may potentially create a large amount of log files, so it is best to do this for initial monitoring purposes to determine the types of websites your users are accessing. After determining the categories that your company approves of, those categories should then be set to allow, which will not generate logs. You can also reduce URL filtering logs by enabling the **Log container page only** option in the URL Filtering profile, so only the main page that matches the category will be logged, not subsequent pages/categories that may be loaded within the container page.

If you subscribe to third-party URL feeds and want to secure your users from emerging threats, see Use an External Dynamic List in a URL Filtering Profile.

STEP 1 | Create a new URL Filtering profile.

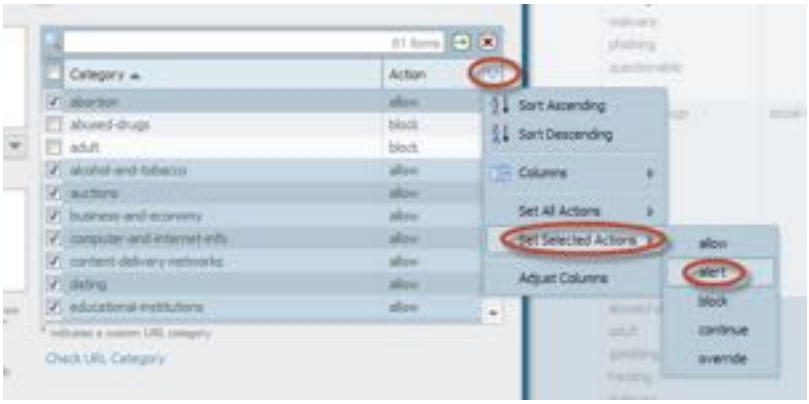
1. Select **Objects** > **Security Profiles URL Filtering**.
2. Select the default profile and then click **Clone**. The new profile will be named **default-1**.
3. Select the **default-1** profile and rename it. For example, rename it to URL-Monitoring.

STEP 2 | Configure the action for all categories to **alert**, except for threat#prone categories, which should remain blocked.



To select all items in the category list from a Windows system, click the first category, then hold down the shift key and click the last category—this will select all categories. Hold the control key (ctrl) down and click items that should be deselected. On a Mac, do the same using the shift and command keys. You could also just set all categories to alert and manually change the recommended categories back to block.

1. In the section that lists all URL categories, select all categories.
2. To the right of the *Action* column heading, mouse over and select the down arrow and then select **Set Selected Actions** and choose **alert**.



3. To ensure that you block access to threat-prone sites, select the following categories and then set the action to **block**: abused-drugs, adult, gambling, hacking, malware. phishing, questionable, weapons.
4. Click **OK** to save the profile.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Apply the URL Filtering profile to the security policy rule(s) that allows web traffic for users.

1. Select **Policies** > **Security** and select the appropriate security policy to modify it.
2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the new profile.
3. Click **OK** to save.

STEP 4 | Save the configuration.

Click **Commit**.

STEP 5 | View the URL filtering logs to determine all of the website categories that your users are accessing. In this example, some categories are set to block, so those categories will also appear in the logs.

For information on viewing the logs and generating reports, see Monitor Web Activity.

Select **Monitor** > **Logs** > **URL Filtering**. A log entry will be created for any website that exists in the URL filtering database that is in a category that is set to any action other than **allow**.

© 2018 Palo Alto Networks, Inc.

Configure URL Filtering

After you Determine URL Filtering Policy Requirements, you should have a basic understanding of what types of websites and website categories your users are accessing. With this information, you are now ready to create custom URL filtering profiles and attach them to the security policy rule(s) that allow web access. In addition to managing web access with a URL Filtering profile, and if you have User-ID configured, you can also manage the sites to which users can submit corporate credentials.

STEP 1 | Create a URL Filtering profile.



If you have not done so already, configure a best practice URL Filtering profile to ensure protection against URLs that have been observed hosting malware or exploitive content.

Select **Objects** > **Security Profiles** > **URL Filtering** and **Add** or modify a URL Filtering profile.

STEP 2 | Define site access for each URL category.

Select **Categories** and set the Site Access for each URL category:

- **Allow** traffic to the URL category. Allowed traffic is not logged.
- Select **alert** to have visibility into sites users are accessing. Matching traffic is allowed, but a URL Filtering log is generated to record when a user accesses a site in the category.
- Select **block** to deny access to traffic that matches the category and to enable logging of the blocked traffic.
- Select **continue** to display a page to users with a warning and require them to click Continue to proceed to a site in the category.
- To only allow access if users provide a configured password, select **override**. For more details on this setting, see Allow Password Access to Certain Sites.

STEP 3 | Configure the URL Filtering profile to detect corporate credential submissions to websites that are in allowed URL categories.



The firewall automatically skips checking credential submissions for App-IDs associated with sites that have never been observed hosting malware or phishing content to ensure the best performance and a low false positive rate even if you enable checks in the corresponding category. The list of sites on which the firewall will skip credential checking is automatically updated via Application and Threat content updates.

1. Select **User Credential Detection**.
2. Select one of the Methods to Check for Corporate Credential Submissions to web pages from the **User Credential Detection** drop-down:
   - **Use IP User Mapping**—Checks for valid corporate username submissions and verifies that the username matches the user logged in the source IP address of the session. To use this method, the firewall matches the submitted username against its IP-address-to-username mapping table. To use this method you can use any of the user mapping methods described in Map IP Addresses to Users.
   - **Use Domain Credential Filter**—Checks for valid corporate usernames and password submissions verifies that the username maps to the IP address of the logged in user. See Configure User Mapping Using the Windows User-ID Agent for instructions on how to set up User-ID to enable this method.
   - **Use Group Mapping**—Checks for valid username submissions based on the user-to-group mapping table populated when you configure the firewall to Map Users to Groups.

© 2018 Palo Alto Networks, Inc.

With group mapping, you can apply credential detection to **any** part of the directory, or specific group, such as groups like IT that have access to your most sensitive applications.



This method is prone to false positives in environments that do not have uniquely structured usernames. Because of this, you should only use this method to protect your high-value user accounts.

3. Set the **Valid Username Detected Log Severity** the firewall uses to log detection of corporate credential submissions. By default, the firewall logs these events as medium severity.

STEP 4 | Allow or block users from submitting corporate credentials to sites based on URL category to Prevent Credential Phishing.



The firewall automatically skips checking credential submissions for App-IDs associated with sites that have never been observed hosting malware or phishing content to ensure the best performance and a low false positive rate even if you enable checks in the corresponding category. The list of sites on which the firewall will skip credential checking is automatically updated via Application and Threat content updates.

1. For each URL category to which **Site Access** is allowed, select how you want to treat **User Credential Submissions**:
   - **alert**—Allow users to submit credentials to the website, but generate a URL Filtering alert log each time a user submits credentials to sites in this URL category.
   - **allow**—(default) Allow users to submit credentials to the website.
   - **block**—Displays the Anti Phishing Block Page to block users from submitting credentials to the website.
   - **continue**—Present the Anti Phishing Continue Page to require users to click Continue to access the site.
2. Configure the URL Filtering profile to detect corporate credential submissions to websites that are in allowed URL categories.

STEP 5 | Define URL Category Exception Lists to specify websites that should always be blocked or allowed, regardless of URL category.

For example, to reduce URL Filtering logs, you may want add you corporate websites in the allow list, so no logs will be generated for those sites. Or, if there is a website this is being overly used and is not work related in any way, you can add it to the block list.

Items in the block list will always be blocked regardless of the action for the associated category, and URLs in the allow list will always be allowed.

For more information on the proper format and wildcards usage, see URL Category Exception Lists.

1. Select **Overrides** and enter URLs or IP addresses in the **Block List** and select an action:
   - **block**—Block the URL.
   - **continue**—Prompt users click **Continue** to proceed to the web page.
   - **override**—The user will be a prompted for a password to continue to the website.
   - **alert**—Allow the user to access the website and add an alert log entry in the URL log.
2. For the **Allow** list, enter IP addresses or URLs that should always be allowed. Each row must be separated by a new line.

STEP 6 | Enable Safe Search Enforcement.

STEP 7 | Log only Container Pages for URL filtering events.

© 2018 Palo Alto Networks, Inc.

1. Select **URL Filtering Settings**. The **Log container page only** option is enabled by default so that only the main page that matches the category is logged, not subsequent pages/categories that may be loaded within the container page.
2. To enable logging for all pages/categories, clear the **Log container page only** check box.

STEP 8 | Enable HTTP Header Logging for one or more of the supported HTTP header fields.

Select **URL Filtering Settings** and select one or more of the following fields to log:

- **User-Agent**
- **Referer**
- **X-Forwarded-For**

STEP 9 | Save the URL Filtering profile and commit your changes.

1. Click **OK**.
2. Click **Commit**.



To test the URL filtering configuration, simply access a website in a category that is set to block or continue to see if the appropriate action is performed.

© 2018 Palo Alto Networks, Inc.

Use an External Dynamic List in a URL Filtering Profile

An External Dynamic List is a text file that is hosted on an external web server. You can use this list to import URLs and enforce policy on these URLs. When the list is updated on the web server, the firewall retrieves the changes and applies policy to the modified list without requiring a commit on the firewall.

For more information, see External Dynamic List.

STEP 1 | Configure the Firewall to Access an External Dynamic List.

- Ensure that the list does not include IP addresses or domain names; the firewall skips non-URL entries.
- Verify the formatting of the list (see Block and Allow Lists).
- Select **URL List** from the Type drop-down.

STEP 2 | Use the external dynamic list in a URL Filtering profile.

1. Select **Objects** > **Security Profiles** > **URL Filtering**.
2. **Add** or modify an existing URL Filtering profile.
3. **Name** the profile and, in the **Categories** tab, select the external dynamic list from the Category list.
4. Click Action to select a more granular action for the URLs in the external dynamic list.

   If a URL that is included in an external dynamic list is also included in a custom URL category, or Block and Allow Lists, the action specified in the custom category or the block and allow list will take precedence over the external dynamic list.
5. Click **OK**.
6. Attach the URL Filtering profile to a Security policy rule.
   1. Select **Policies** > **Security**.
   2. Select the **Actions** tab and, in the Profile Setting section, select the new profile in the **URL Filtering** drop-down.
   3. Click **OK** and **Commit**.

STEP 3 | Test that the policy action is enforced.

1. View External Dynamic List Entries for the URL list, and attempt to access a URL from the list.
2. Verify that the action you defined is enforced in the browser.
3. To monitor the activity on the firewall:
   1. Select **ACC** and add a URL Domain as a global filter to view the Network Activity and Blocked Activity for the URL you accessed.
   2. Select **Monitor** > **Logs** > **URL Filtering** to access the detailed log view.

STEP 4 | Verify whether entries in the external dynamic list were ignored or skipped.

In a list of type URL, the firewall skips non-URL entries as invalid and ignores entries that exceed the maximum limit for the firewall model.

To check whether you have reached the limit for an external dynamic list type, select Objects > External Dynamic Lists and click List Capacities.

Use the following CLI command on a firewall to review the details for a list.

© 2018 Palo Alto Networks, Inc.

```
request system external-list show type url name <list_name>
```

For example:

```
request system external-list show type url name My_URL_List
vsys5/My_URL_List:
Next update at: Tue Jan 3 14:00:00 2017
Source: http://example.com/My_URL_List.txt
Referenced: Yes
Valid: Yes
Auth-Valid: Yes

Total valid entries: 3
Total invalid entries: 0
Valid urls:
www.URL1.com
www.URL2.com
www.URL3.com
```


© 2018 Palo Alto Networks, Inc.

Customize the URL Filtering Response Pages

The firewall provides predefined URL Filtering Response Pages on page 652 that display by default when a user:

- A user attempts to browse to a site in a category with restricted access.
- A user submits valid corporate credentials to a site for which credential detection is enabled (Prevent Credential Phishing on page 557 based on URL category).
- Container Pages on page 651 blocks a search attempt.

However, you can create your own custom response pages with your corporate branding, acceptable use policies, and links to your internal resources.



Custom response pages larger than the maximum supported size are not decrypted or displayed to users. In PAN-OS 8.0.11 and earlier releases, custom response pages on a decrypted site cannot exceed 8,191 bytes; the maximum size is increased to 17,999 bytes in PAN-OS 8.0.12 and later PAN-OS 8.0 releases.

STEP 1 | Export the default response page(s).

1. Select **Device** > **Response Pages**.
2. Select the link for the URL filtering response page you want to modify.
3. Click the response page (predefined or shared) and then click the **Export** link and save the file to your desktop.

STEP 2 | Edit the exported page.

1. Using the HTML text editor of your choice, edit the page:
   - If you want the response page to display custom information about the specific user, URL, or category that was blocked, add one or more of the supported Table 4: URL Filtering Response Page Variables on page 654.
   - If you want to include custom images (such as your corporate logo), a sound, or style sheet, or link to another URL, for example to a document detailing your acceptable web use policy, include one or more of the supported Table 5: Response Page References on page 655.
2. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding. For example, in Notepad you would select **UTF-8** from the **Encoding** drop-down in the Save As dialog.

STEP 3 | Import the customized response page.

1. Select **Device** > **Response Pages**.
2. Select the link that corresponds to the URL Filtering response page you edited.
3. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.
4. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select **shared** to make it available to all virtual systems.
5. Click **OK** to import the file.

STEP 4 | Save the new response page(s).

Commit the changes.

STEP 5 | Verify that the new response page displays.

From a browser, go to the URL that will trigger the response page. For example, to see a modified URL Filtering and Category Match response page, browse to URL that your URL filtering policy is set to block.

© 2018 Palo Alto Networks, Inc.

Allow Password Access to Certain Sites

In some cases there may be URL categories that you want to block, but allow certain individuals to browse to on occasion. In this case, you would set the category action to **override** and define a URL admin override password in the firewall Content-ID configuration. When users attempt to browse to the category, they will be required to provide the override password before they are allowed access to the site. Use the following procedure to configure URL admin override:

STEP 1 | Set the URL admin override password.

1. Select **Device** > **Setup** > **Content ID**.
2. In the **URL Admin Override** section, click **Add**.
3. In the **Location** field, select the virtual system to which this password applies.
4. Enter the **Password** and **Confirm Password**.
5. Select an **SSL/TLS Service Profile**. The profile specifies the certificate that the firewall presents to the user if the site with the override is an HTTPS site. For details, see Configure an SSL/TLS Service Profile.
6. Select the **Mode** for prompting the user for the password:
   - **Transparent**—The firewall intercepts the browser traffic destined for site in a URL category you have set to override and impersonates the original destination URL, issuing an HTTP 401 to prompt for the password.

     The client browser will display certificate errors if it does not trust the certificate.
   - **Redirect**—The firewall intercepts HTTP or HTTPS traffic to a URL category set to override and redirects the request to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to prompt for the override password. If you select this option, you must provide the **Address** (IP address or DNS hostname) to which to redirect the traffic.
7. Click **OK**.

STEP 2 | (Optional) Set a custom override period.

1. Edit the URL Filtering section.
2. To change the amount of time users can browse to a site in a category for which they have successfully entered the override password, enter a new value in the **URL Admin Override Timeout** field. By default, users can access sites within the category for 15 minutes without re-entering the password.
3. To change the amount of time users are blocked from accessing a site set to override after three failed attempts to enter the override password, enter a new value in the **URL Admin Lockout Timeout** field. By default, users are blocked for 30 minutes.
4. Click **OK**.

STEP 3 | (Redirect mode only) Create a Layer 3 interface to which to redirect web requests to sites in a category configured for override.

1. Create a management profile to enable the interface to display the URL Filtering Continue and Override Page response page:
   1. Select **Network** > **Interface Mgmt** and click **Add**.
   2. Enter a **Name** for the profile, select **Response Pages**, and then click **OK**.
2. Create the Layer 3 interface. Be sure to attach the management profile you just created (on the **Advanced** > **Other Info** tab of the Ethernet Interface dialog).

© 2018 Palo Alto Networks, Inc.

STEP 4 | (Redirect mode only) To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface to which you are redirecting web requests to a site in a URL category configured for override.You can either generate a self-signed certificate or import a certificate that is signed by an external CA.

To use a self-signed certificate, you must first create a root CA certificate and then use that CA to sign the certificate you will use for URL admin override as follows:

1. To create a root CA certificate, select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**. Enter a **Certificate Name**, such as RootCA. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed). Make sure you select the **Certificate Authority** check box and then click **Generate** the certificate.
2. To create the certificate to use for URL admin override, click **Generate**. Enter a **Certificate Name** and enter the DNS hostname or IP address of the interface as the **Common Name**. In the **Signed By** field, select the CA you created in the previous step. Add an IP address attribute and specify the IP address of the Layer 3 interface to which you will be redirecting web requests to URL categories that have the override action.
3. **Generate** the certificate.
4. To configure clients to trust the certificate, select the CA certificate on the **Device Certificates** tab and click **Export**. You must then import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO).

STEP 5 | Specify which URL categories require an override password to enable access.

1. Select **Objects** > **URL Filtering** and either select an existing URL filtering profile or **Add** a new one.
2. On the **Categories** tab, set the Action to **override** for each category that requires a password.
3. Complete any remaining sections on the URL filtering profile and then click **OK** to save the profile.

STEP 6 | Apply the URL Filtering profile to the security policy rule(s) that allows access to the sites requiring password override for access.

1. Select **Policies** > **Security** and select the appropriate security policy to modify it.
2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the profile.
3. Click **OK** to save.

STEP 7 | Save the configuration.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos in search query return traffic. You can enable the firewall to block search results if the end user is not using the strictest safe search settings, and you can also transparently enable safe search for your users. The firewall supports safe search enforcement for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. Consider that safe search is a best-effort setting and service providers do not guarantee that it works with every website, and search providers classify sites as safe or unsafe (not Palo Alto Networks).

To use this feature you must enable the **Safe Search Enforcement** option in a URL filtering profile and attach it to a security policy rule.The firewall then blocks any matching search query return traffic that is not using the strictest safe search settings. There are two methods to enforce safe search:

- Block Search Results when Strict Safe Search is not Enabled—When an end user attempts to perform a search without first enabling the strictest safe search settings, the firewall blocks the search query results and displays the URL Filtering Safe Search Block Page. By default, this page will provide a URL to the search provider settings for configuring safe search.
- Transparently Enable Safe Search for Users—When an end user attempts to perform a search without first enabling the strict safe search settings, the firewall blocks the search results with an HTTP 503 status code and redirects the search query to a URL that includes the safe search parameters. You enable this functionality by importing a new URL Filtering Safe Search Block Page containing the JavaScript for rewriting the search URL to include the strict safe search parameters. In this configuration, users will not see the block page, but will instead be automatically redirected to a search query that enforces the strictest safe search options. This safe search enforcement method requires content release version 475 or later and is only supported for Google, Yahoo, and Bing searches.

As safe search settings differ by search provider, get started by reviewing the different safe search implementations. There are then two ways you can enforce safe search: you can block search results when safe search is disabled, or you can transparently enable safe search for your users:

- Safe Search Settings for Search Providers
- Block Search Results when Strict Safe Search is not Enabled
- Transparently Enable Safe Search for Users

Safe Search Settings for Search Providers

Safe search settings differ for each search provider—review the following settings to learn more.

Search Provider	Safe Search Setting Description
Google/YouTube	Offers safe search on individual computers or network-wide through Google's safe search virtual IP address: **Safe Search Enforcement for Google Searches on Individual Computers** In the Google Search Settings, the **Filter explicit results** setting enables safe search functionality. When enabled, the setting is stored in a browser cookie as `FF=` and passed to the server each time the user performs a Google search. Appending **`safe=active`** to a Google search query URL also enables the strictest safe search settings. **Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address**

© 2018 Palo Alto Networks, Inc.

Search Provider	Safe Search Setting Description
	Google provides servers that Lock SafeSearch (forcesafesearch.google.com) settings in every Google and YouTube search. By adding a DNS entry for **`www.google.com`** and **`www.youtube.com`** (and other relevant Google and YouTube country subdomains) that includes a CNAME record pointing to **`forcesafesearch.google.com`** to your DNS server configuration, you can ensure that all users on your network are using strict safe search settings every time they perform a Google or YouTube search. Keep in mind, however, that this solution is not compatible with Safe Search Enforcement on the firewall. Therefore, if you are using this option to force safe search on Google, the best practice is to block access to other search engines on the firewall by creating custom URL categories and adding them to the block list in the URL filtering profile. *If you plan to use the Google Lock SafeSearch solution, consider configuring DNS Proxy (Network > DNS Proxy) and setting the inheritance source as the Layer 3 interface on which the firewall receives DNS settings from service provider via DHCP. You would configure the DNS proxy with Static Entries for www.google.com and www.youtube.com, using the local IP address for the forcesafesearch.google.com server.*
Yahoo	Offers safe search on individual computers only. The Yahoo Search Preferences includes three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `vm=` and passed to the server each time the user performs a Yahoo search. Appending `vm=r` to a Yahoo search query URL also enables the strictest safe search settings. *When performing a search on Yahoo Japan (**`yahoo.co.jp`**) while logged into a Yahoo account, end users must also enable the SafeSearch Lock option.*
Bing	Offers safe search on individual computers or through their Bing in the Classroom program. The Bing Settings include three SafeSearch settings: **Strict**, **Moderate**, or **Off**. When enabled, the setting is stored in a browser cookie as `adtl=` and passed to the server each time the user performs a Bing search. Appending **`adlt=strict`** to a Bing search query URL also enables the strictest safe search settings. The Bing SSL search engine does not enforce the safe search URL parameters and you should therefore consider blocking Bing over SSL for full safe search enforcement.

Block Search Results when Strict Safe Search is not Enabled

By default, when you enable safe search enforcement, when a user attempts to perform a search without using the strictest safe search settings, the firewall will block the search query results and display the URL Filtering Safe Search Block Page. This page provides a link to the search settings page for the corresponding search provider so that the end user can enable the safe search settings. If you plan to use this default method for enforcing safe search, you should communicate the policy to your end users prior to deploying the policy. See for details on how each search provider implements safe search. The default URL Filtering

© 2018 Palo Alto Networks, Inc.

Safe Search Block Page provides a link to the search settings for the corresponding search provider. You can optionally Customize the URL Filtering Response Pages.

Alternatively, to enable safe search enforcement so that it is transparent to your end users, configure the firewall to Transparently Enable Safe Search for Users.

STEP 1 | Enable Safe Search Enforcement in the URL Filtering profile.

1. Select **Objects** > **Security Profiles** > **URL Filtering**.
2. Select an existing profile to modify, or clone the default profile to create a new profile.
3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.
4. (Optional) Restrict users to specific search engines:
    1. On the **Categories** tab, set the **search-engines** category to **block**.
    2. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:

        www.google.com

        www.bing.com
5. Configure other settings as necessary to:
    - Define site access for each URL category.
    - Define Block and Allow Lists to specify websites that should always be blocked or allowed, regardless of URL category.
6. Click **OK** to save the profile.

STEP 2 | Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet.

1. Select **Policies** > **Security** and select a rule to which to apply the URL filtering profile that you just enabled for Safe Search Enforcement.
2. On the **Actions** tab, select the **URL Filtering** profile.
3. Click **OK** to save the security policy rule.

STEP 3 | Enable SSL Forward Proxy decryption.

Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings.

1. Add a custom URL category for the search sites:
    1. Select **Objects** > **Custom Objects** > **URL Category** and **Add** a custom category.
    2. Enter a **Name** for the category, such as SearchEngineDecryption.
    3. **Add** the following to the Sites list:

        www.bing.*

        www.google.*

        search.yahoo.*
    4. Click **OK** to save the custom URL category object.
2. Follow the steps to Configure SSL Forward Proxy.
3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**.

STEP 4 | (Recommended) Block Bing search traffic running over SSL.

© 2018 Palo Alto Networks, Inc.

Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL.

1. Add a custom URL category for Bing:
    1. Select **Objects** > **Custom Objects** > **URL Category** and **Add** a custom category.
    2. Enter a **Name** for the category, such as EnableBingSafeSearch.
    3. **Add** the following to the Sites list:

        www.bing.com/images/*

        www.bing.com/videos/*
    4. Click **OK** to save the custom URL category object.
2. Create another URL filtering profile to block the custom category you just created:
    1. Select **Objects** > **Security Profiles** > **URL Filtering**.
    2. **Add** a new profile and give it a descriptive **Name**.
    3. Locate the custom category in the Category list and set it to **block**.
    4. Click **OK** to save the URL filtering profile.
3. Add a security policy rule to block Bing SSL traffic:
    1. Select **Policies** > **Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.
    2. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.
    3. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.
    4. Select **TCP** as the **Protocol** and set the **Destination Port** to `443`.
    5. Click **OK** to save the rule.
    6. Use the **Move** options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled.

STEP 5 | Save the configuration.

Click **Commit**.

STEP 6 | Verify the Safe Search Enforcement configuration.

This verification step only works if you are using block pages to enforce safe search. If you are using transparent safe search enforcement, the firewall block page will invoke a URL rewrite with the safe search parameters in the query string.

1. From a computer that is behind the firewall, disable the strict search settings for one of the supported search providers. For example, on bing.com, click the **Preferences** icon on the Bing menu bar.

    
2. Set the **SafeSearch** option to **Moderate** or **Off** and click **Save**.
3. Perform a Bing search and verify that the URL Filtering Safe Search Block page displays instead of the search results:

© 2018 Palo Alto Networks, Inc.



4. Use the link in the block page to go to the search settings for the search provider and set the safe search setting back to the strictest setting (**Strict** in the case of Bing) and then click **Save**.
5. Perform a search again from Bing and verify that the filtered search results display instead of the block page.

Transparently Enable Safe Search for Users

If you want to enforce filtering of search query results with the strictest safe search filters, but you don't want your end users to have to manually configure the settings, you can enable transparent safe search enforcement as follows. This functionality is supported on Google, Yahoo, and Bing search engines only and requires Content Release version 475 or later.

STEP 1 | Make sure the firewall is running Content Release version 475 or later.

1. Select **Device** > **Dynamic Updates**.
2. Check the **Applications and Threats** section to determine what update is currently running.
3. If the firewall is not running the required update or later, click **Check Now** to retrieve a list of available updates.
4. Locate the required update and click **Download**.
5. After the download completes, click **Install**.

STEP 2 | Enable Safe Search Enforcement in the URL Filtering profile.

1. Select **Objects** > **Security Profiles** > **URL Filtering**.
2. Select an existing profile to modify, or clone the default profile to create a new one.
3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.
4. (Optional) Allow access to specific search engines only:
    1. On the **Categories** tab, set the **search-engines** category to **block**.
    2. For each search engine that you want end users to be able to access, enter the web address in the **Allow List** text box. For example, to allow users access to Google and Bing searches only, you would enter the following:

        `www.google.com`

        `www.bing.com`
5. Configure other settings as necessary to:
    - Define site access for each URL category.
    - Define Block and Allow Lists to specify websites that should always be blocked or allowed, regardless of URL category.
6. Click **OK** to save the profile.

STEP 3 | Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet.

1. Select **Policies** > **Security** and select a rule to which to apply the URL filtering profile that you just enabled for Safe Search Enforcement.
2. On the **Actions** tab, select the **URL Filtering** profile.

© 2018 Palo Alto Networks, Inc.

3. Click **OK** to save the security policy rule.

STEP 4 | (Recommended) Block Bing search traffic running over SSL.

Because the Bing SSL search engine does not adhere to the safe search settings, for full safe search enforcement, you must deny all Bing sessions that run over SSL.

1. Add a custom URL category for Bing:
    1. Select **Objects** > **Custom Objects** > **URL Category** and **Add** a custom category.
    2. Enter a **Name** for the category, such as EnableBingSafeSearch.
    3. **Add** the following to the Sites list:

        www.bing.com/images/*

        www.bing.com/videos/*
    4. Click **OK** to save the custom URL category object.
2. Create another URL filtering profile to block the custom category you just created:
    1. Select **Objects** > **Security Profiles** > **URL Filtering**.
    2. **Add** a new profile and give it a descriptive **Name**.
    3. Locate the custom category you just created in the Category list and set it to **block**.
    4. Click **OK** to save the URL filtering profile.
3. **Add** a security policy rule to block Bing SSL traffic:
    1. Select **Policies** > **Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.
    2. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.
    3. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.
    4. Select **TCP** as the **Protocol**, set the **Destination Port** to **443**.
    5. Click **OK** to save the rule.
    6. Use the **Move** options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled.

STEP 5 | Edit the URL Filtering Safe Search Block Page, replacing the existing code with the JavaScript for rewriting search query URLs to enforce safe search transparently.

1. Select **Device** > **Response Pages** > **URL Filtering Safe Search Block Page**.
2. Select **Predefined** and then click **Export** to save the file locally.
3. Use an HTML editor and replace all of the existing block page text with the following text and then save the file.

```
<html>
  <head>
    <title>Search Blocked</title>
    <meta http-equiv="Content-Type" content="text/html; charset=utf-8">
    <meta http-equiv="pragma" content="no-cache">
    <meta name="viewport" content="initial-scale=1.0">
    <style>
      #content {
      border:3px solid#aaa;
      background-color:#fff;
      margin:1.5em;
      padding:1.5em;
      font-family:Tahoma,Helvetica,Arial,sans-serif;
      font-size:1em;
      }
```


© 2018 Palo Alto Networks, Inc.

```
      h1 {
      font-size:1.3em;
      font-weight:bold;
      color:#196390;
      }
      b {
      font-weight:normal;
      color:#196390;
      }
    </style>
  </head>
  <body bgcolor="#e7e8e9">
    <div id="content">
      <h1>Search Blocked</h1>
      <p>
        <b>User:</b>
        <user/>
      </p>
      <p>Your search results have been blocked because your search
 settings are not in accordance with company policy.  In order to
 continue, please update your search settings so that Safe Search is
 set to the strictest setting. If you are currently logged into your
 account, please also lock Safe Search and try your search again.</p>
      <p>
        For more information, please refer to:
        <a href="<ssurl/>">
          <ssurl/>
        </a>
      </p>
      <p id="java_off"> Please enable JavaScript in your browser.<br></
p>
      <p><b>Please contact your system administrator if you believe this
 message is in error.</b></p>
    </div>
  </body>
  <script>
    // Grab the URL that's in the browser.
    var s_u = location.href;
    //bing
    // Matches the forward slashes in the beginning, anything, then
 ".bing." then anything followed by a non greedy slash. Hopefully the
 first forward slash.
    var b_a = /^.*\/\/(.+\.bing\..+?)\//.exec(s_u);
    if (b_a) {
       s_u = s_u + "&adlt=strict";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You are being
 redirected to a safer search!';
    }
    //google
    // Matches the forward slashes in the beginning, anything, then
 ".google." then anything followed by a non greedy slash. Hopefully the
 first forward slash.
    var g_a = /^.*\/\/(.+\.google\..+?)\//.exec(s_u);
    if (g_a) {
        s_u = s_u.replace(/&safe=off/ig,"");
        s_u = s_u + "&safe=active";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You are being
 redirected to a safer search!';    }
    //yahoo
```


© 2018 Palo Alto Networks, Inc.

```
    // Matches the forward slashes in the beginning, anything, then
 ".yahoo."" then anything followed by a non greedy slash. Hopefully the
 first forward slash.
    var y_a = /^.*\/\/(.+\.yahoo\..+?)\//.exec(s_u);
    if (y_a) {
        s_u = s_u.replace(/&vm=p/ig,"");
        s_u = s_u + "&vm=r";
        window.location.replace(s_u);
        document.getElementById("java_off").innerHTML = 'You are being
 redirected to a safer search!';
    }
    document.getElementById("java_off").innerHTML = ' ';
  </script>
</html>
```

STEP 6 | Import the edited URL Filtering Safe Search Block page onto the firewall.

1. To import the edited block page, select **Device** > **Response Pages** > **URL Filtering Safe Search Block Page**.
2. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file.
3. (Optional) Select the virtual system on which this login page will be used from the **Destination** drop-down or select **shared** to make it available to all virtual systems.
4. Click **OK** to import the file.

STEP 7 | Enable SSL Forward Proxy decryption.

Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings.

1. Add a custom URL category for the search sites:
    1. Select **Objects** > **Custom Objects** > **URL Category** and **Add** a custom category.
    2. Enter a **Name** for the category, such as SearchEngineDecryption.
    3. **Add** the following to the Sites list:

        ```
        www.bing.*

        www.google.*

        search.yahoo.*
        ```
    4. Click **OK** to save the custom URL category object.
2. Follow the steps to Configure SSL Forward Proxy.
3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**.

STEP 8 | Save the configuration.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Monitor Web Activity

The ACC, URL filtering logs and reports show all user web activity for URL categories that are set to **alert**, **block**, **continue**, or **override**. By monitoring the logs, you can gain a better understanding of the web activity of your user base to determine a web access policy.

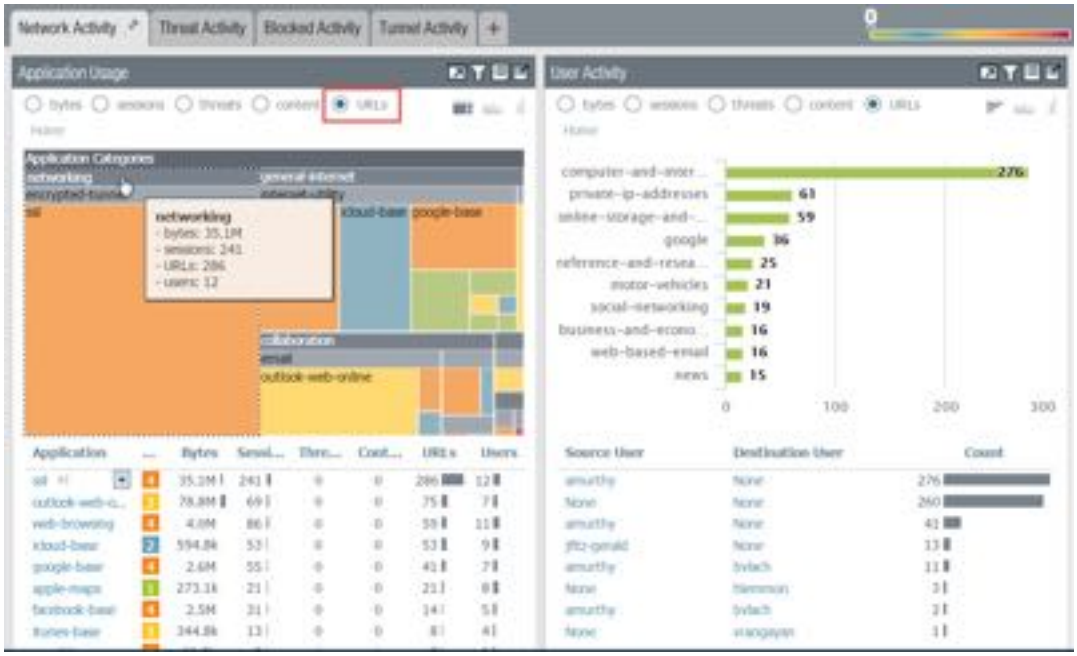
The following topics describe how to monitor web activity:

- Monitor Web Activity of Network Users on page 680
- View the User Activity Report on page 682
- Configure Custom URL Filtering Reports on page 683

Monitor Web Activity of Network Users

You can use the ACC, URL filtering reports, and logs that are generated on the firewall to track user activity.

- For a quick view of the most common categories users access in your environment, check the **ACC** widgets. Most **Network Activity** widgets allow you to sort on URLs. For example, in the Application Usage widget, you can see that the networking category is the most accessed category, followed by encrypted tunnel, and ssl. You can also view the list of **Threat Activity** and **Blocked Activity** sorted on URLs.



  View logs and configure log options:

- From the ACC, you can jump directly to the logs (■) or select **Monitor** > **Logs** > **URL Filtering**.

  The log action for each entry depends on the Site Access setting you defined for the corresponding category:

  - **Alert log**—In this example, the computer-and-internet-info category is set to alert.

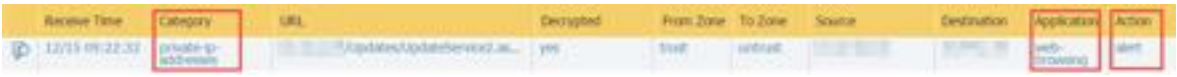


© 2018 Palo Alto Networks, Inc.

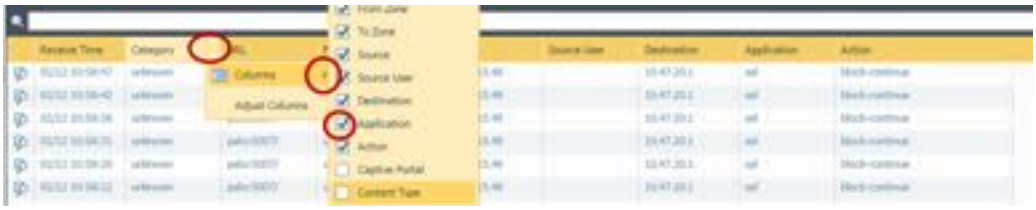
- **Block log**—In this example, the insufficient-content category is set to continue. If the category had been set to block instead, the log Action would be block-url.



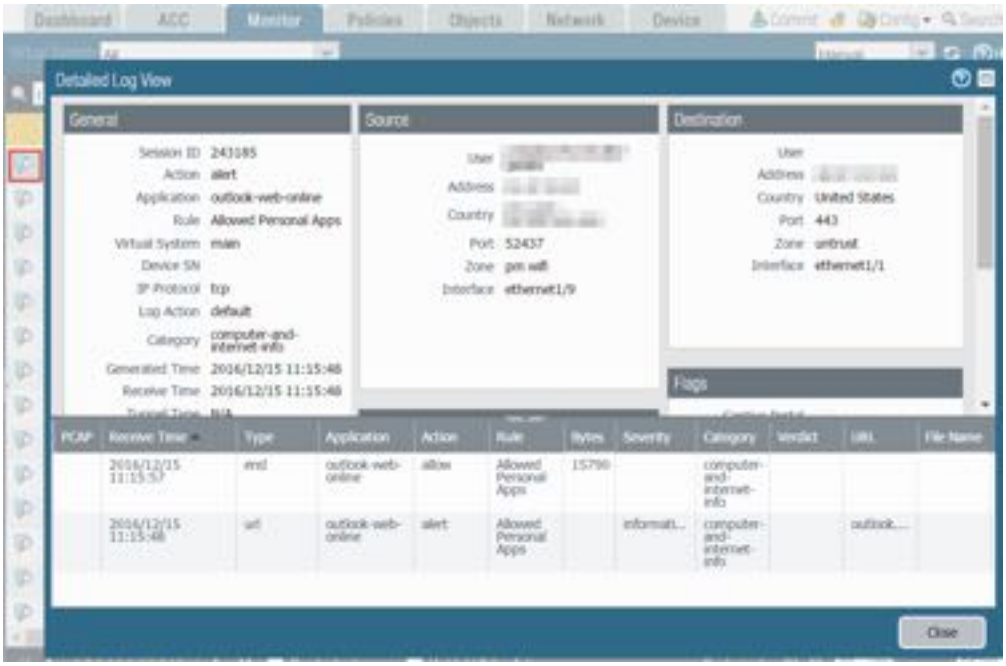
- **Alert log on encrypted website**—In this example, the category is private-ip-addresses and the application is web-browsing. This log also indicates that the firewall decrypted this traffic.



- You can also add several other columns to your URL Filtering log view, such as: to and from zone, content type, and whether or not a packet capture was performed. To modify what columns to display, click the down arrow in any column and select the attribute to display.



- To view the complete log details and/or request a category change for the given URL that was accessed, click the log details icon in the first column of the log.



- Generate predefined URL filtering reports on URL categories, URL users, Websites accessed, Blocked categories, and more.

  Select **Monitor** > **Reports** and under the **URL Filtering Reports** section, select one of the reports. The reports cover the 24-hour period of the date you select on the calendar. You can also export the report to PDF, CSV, or XML.

© 2018 Palo Alto Networks, Inc.



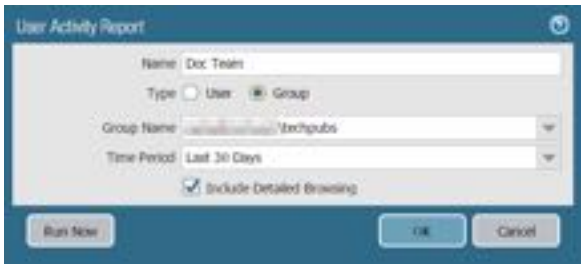
View the User Activity Report

This report provides a quick method of viewing user or group activity and also provides an option to view browse time activity.

STEP 1 | Configure a User Activity Report.

1. Select **Monitor** > **PDF Reports** > **User Activity Report**.
2. **Add** a report and enter a **Name** for it.
3. Select the report **Type**:
   - Select **User** to generate a report for one person.
   - Select **Group** for a group of users.

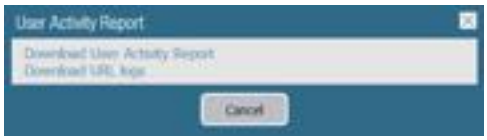
   You must Enable User-ID in order to be able to select user or group names. If User-ID is not configured, you can select the type User and enter the IP address of the user's computer.
4. Enter the **Username/IP Address** for a user report or enter the group name for a user group report.
5. Select the time period. You can select an existing time period, or select **Custom**.
6. Select the **Include Detailed Browsing** check box, so browsing information is included in the report.



STEP 2 | Run the report.

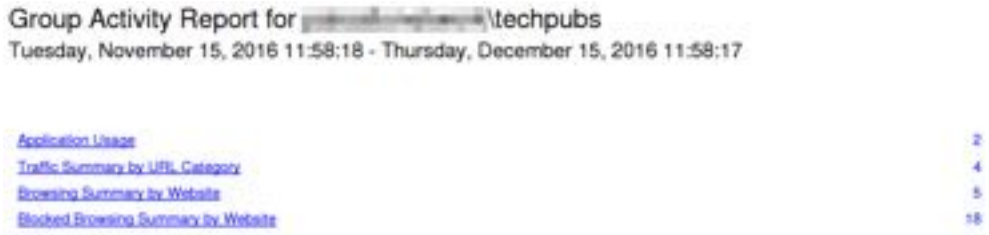
© 2018 Palo Alto Networks, Inc.

1. Click **Run Now**.
2. When the firewall finishes generating report, click one of the links to download it:
   - Click **Download User Activity Report** to download a PDF version of the report.
   - Click **Download URL Logs** to download a CSV file of the corresponding log entries.

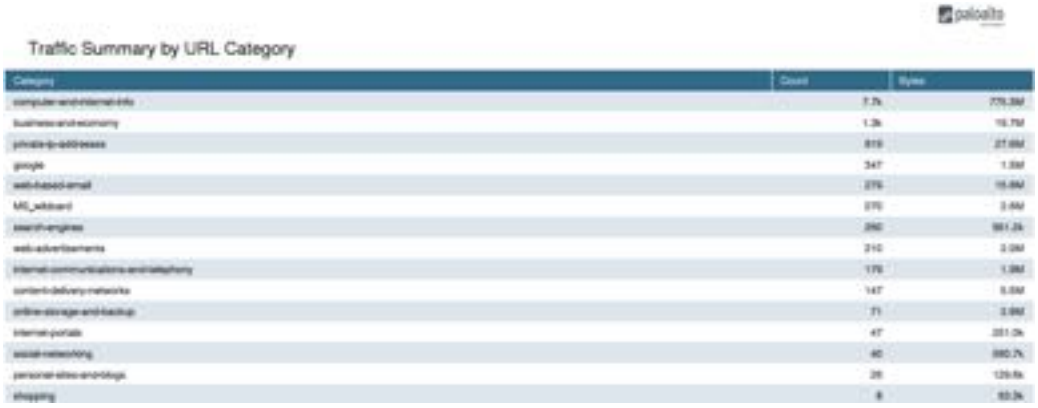


3. After downloading the report, click **Cancel**.
4. If you want to save the user activity report settings so you can run the same report again later, click **OK**; otherwise click **Cancel**.

STEP 3 | View the user activity report by opening the file that you downloaded. The PDF version of the report shows the user or group on which you based the report, the report time frame, and a table of contents:



STEP 4 | Click an item in the table of contents to view the report details. For example, click **Traffic Summary by URL Category** to view statistics for the selected user or group.



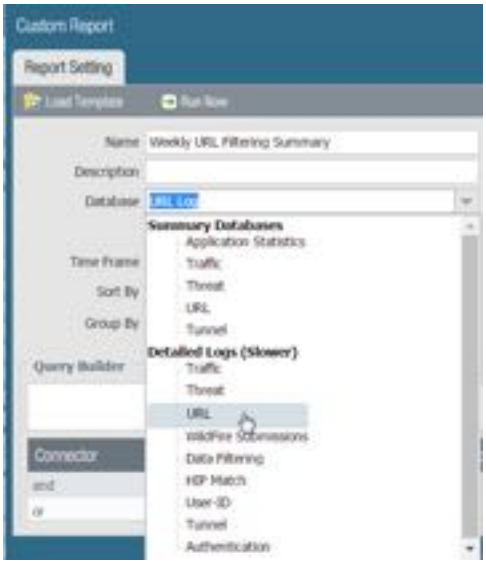
Configure Custom URL Filtering Reports

To generate a detailed report that you can schedule to run regularly, configure a custom URL Filtering report. You can choose any combination of URL Filtering log fields on which to base the report.

STEP 1 | Add a new custom report.

1. Select **Monitor** > **Manage Custom Reports** and **Add** a report.
2. Give the report a unique **Name**, and optionally a **Description**.
3. Select the **Database** you want to use to generate the report. To generate a detailed URL Filtering report, select **URL** from the Detailed Logs section:

© 2018 Palo Alto Networks, Inc.



STEP 2 | Configure report options.

1. Select a predefined **Time Frame** or select **Custom**.
2. Select the log columns to include in the report from the Available Columns list add them (⊕) to the Selected Columns. For example, for a URL Filtering report you might select:
   - Action
   - App Category
   - Category
   - Destination Country
   - Source User
   - URL



3. If the firewall is enabled to Prevent Credential Phishing, select the Attribute **Flags**, the Operator **has** and the Value **Credential Detected** to also include events in the report that record when a user submitted a valid corporate credential to a site.



4. (Optional) Select a **Sort By** option to set the attribute to use to aggregate the report details. If you do not select an attribute to sort by, the report will return the first N number of results without any aggregation. Select a **Group By** attribute to use as an anchor for grouping data. The following example shows a report with **Group By** set to **App Category** and **Sort By** set to a **Count** of **Top 5**.

© 2018 Palo Alto Networks, Inc.



STEP 3 | Run the report.

1. Click the **Run Now** icon to immediately generate the report, which opens in a new tab.
2. When you are done reviewing the report, go back to the **Report Setting** tab and either tune the settings and run the report again, or continue to the next step to schedule the report.
3. Select the **Schedule** check box to run the report once per day. This will generate a daily report that details web activity over the last 24 hours.

STEP 4 | **Commit** the configuration.

STEP 5 | View the custom report.

1. Select **Monitor** > **Reports**.
2. Expand the **Custom Reports** pane in the right column and select the report you want to view. The latest report displays automatically.
3. To view the report for a previous date, select the date from the calendar. You can also export the report to PDF, CSV, or XML.

© 2018 Palo Alto Networks, Inc.

Set Up the PAN-DB Private Cloud

To deploy one or more M-500 appliances as a PAN-DB private cloud within your network or data center, you must complete the following tasks:

- Configure the PAN-DB Private Cloud on page 686
- Configure the Firewalls to Access the PAN-DB Private Cloud on page 689

Configure the PAN-DB Private Cloud

STEP 1 | Rack mount the M-500 appliance.

Refer to the M-500 Hardware Reference Guide for instructions.

STEP 2 | Register the M-500 appliance.

For instructions on registering the M-500 appliance, see Register the Firewall.

STEP 3 | Perform Initial Configuration of the M-500 Appliance.

The M-500 appliance in PAN-DB mode uses two ports- MGT (Eth0) and Eth1; Eth2 is not used in PAN-DB mode. The management port is used for administrative access to the appliance and for obtaining the latest content updates from the PAN-DB public cloud. For communication between the appliance (PAN-DB server) and the firewalls on the network, you can use the MGT port or Eth1.

1. Connect to the M-500 appliance in one of the following ways:
   - Attach a serial cable from a computer to the Console port on the M-500 appliance and connect using a terminal emulation software (9600-8-N-1).
   - Attach an RJ-45 Ethernet cable from a computer to the MGT port on the M-500 appliance. From a browser, go to https://192.168.1.1. Enabling access to this URL might require changing the IP address on the computer to an address in the 192.168.1.0 network (for example, 192.168.1.2).
2. When prompted, log in to the appliance. Log in using the default username and password (admin/admin). The appliance will begin to initialize.
3. Configure network access settings including the IP address for the MGT interface:

```
set deviceconfig system ip-address  <server-IP> netmask <netmask> default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>
```

   where *<server-IP>* is the IP address you want to assign to the management interface of the server, *<netmask>* is the subnet mask, *<gateway-IP>* is the IP address of the network gateway, and *<DNS-IP>* is the IP address of the primary DNS server.
4. Configure network access settings including the IP address for the Eth1 interface:

```
set deviceconfig system eth1 ip-address  <server-IP> netmask <netmask>
 default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>
```

   where *<server-IP>* is the IP address you want to assign to the data interface of the server, *<netmask>* is the subnet mask, *<gateway-IP>* is the IP address of the network gateway, and *<DNS-IP>* is the IP address of the DNS server.
5. Save your changes to the PAN-DB server.

© 2018 Palo Alto Networks, Inc.

commit

STEP 4 | Switch to PAN-DB private cloud mode.

1. To switch to PAN-DB mode, use the CLI command:

```
request system system-mode pan-url-db
```



You can switch from Panorama mode to PAN-DB mode and back; and from Panorama mode to Log Collector mode and back. Switching directly from PAN-DB mode to Log Collector mode or vice versa is not supported. When switching operational mode, a data reset is triggered. With the exception of management access settings, all existing configuration and logs will be deleted on restart.

2. Use the following command to verify that the mode is changed:

```
show pan-url-cloud-status
hostname: M-500
ip-address: 1.2.3.4
netmask: 255.255.255.0
default-gateway: 1.2.3.1
ipv6-address: unknown
ipv6-link-local-address: fe80:00/64
ipv6-default-gateway:
mac-address: 00:56:90:e7:f6:8e
time: Mon Apr 27 13:43:59 2015
uptime: 10 days, 1:51:28
family: m
model: M-500
serial: 0073010000xxx
sw-version: 7.0.0
app-version: 492-2638
app-release-date: 2015/03/19  20:05:33
av-version: 0
av-release-date: unknown
wf-private-version: 0
wf-private-release-date: unknown
logdb-version: 7.0.9
platform-family: m
pan-url-db: 20150417-220
system-mode: Pan-URL-DB
operational-mode: normal
```

3. Use the following command to check the version of the cloud database on the appliance:

```
show pan-url-cloud-status
Cloud status:            Up
URL database version:    20150417-220
```

STEP 5 | Install content and database updates.



The appliance only stores the currently running version of the content and one earlier version.

Pick one of the following methods of installing the content and database updates:

- If the PAN-DB server has direct Internet access use the following commands:

© 2018 Palo Alto Networks, Inc.

1. To check whether a new version is published use:

   request pan-url-db upgrade check
2. To check the version that is currently installed on your server use:

   request pan-url-db upgrade info
3. To download and install the latest version:
   - **request pan-url-db upgrade download latest**
   - **request pan-url-db upgrade install** *<version latest | file>*
4. To schedule the M-500 appliance to automatically check for updates:

```
set deviceconfig system update-schedule pan-url-db recurring weekly
 action download-and-install day-of-week  <day of week> at <hr:min>
```

- If the PAN-DB server is offline, access the Palo Alto Networks Customer Support web site to download and save the content updates to an SCP server on your network. You can then import and install the updates using the following commands:
  - **scp import pan-url-db remote-port** *<port-number>* from username@host:path
  - **request pan-url-db upgrade install file** *<filename>*

STEP 6 | Set up administrative access to the PAN-DB private cloud.



The appliance has a default `admin` *account. Any additional administrative users that you create can either be superusers (with full access) or superusers with read-only access.*

PAN-DB private cloud does not support the use of RADIUS VSAs. If the VSAs used on the firewall or Panorama are used for enabling access to the PAN-DB private cloud, an authentication failure will occur.

- To set up a local administrative user on the PAN-DB server:
  1. **configure**
  2. **set mgt-config users** *<username>* permissions role-based <superreader | superuser> yes
  3. **set mgt-config users** *<username>* password
  4. **Enter password:xxxxx**
  5. **Confirm password:xxxxx**
  6. **commit**
- To set up an administrative user with RADIUS authentication:
  1. Create RADIUS server profile.

```
set shared server-profile radius  <server_profile_name>
 server <server_name> ip-address <ip_address> port <port_no>
 secret <shared_password>
```

  2. Create authentication-profile.

© 2018 Palo Alto Networks, Inc.

```
set shared authentication-profile  <auth_profile_name> user-
domain <domain_name_for_authentication> allow-list <all> method radius
 server-profile <server_profile_name>
```

3. Attach the authentication-profile to the user.

```
set mgt-config users  <username> authentication-
profile <auth_profile_name>
```

4. Commit the changes.

```
commit
```

- To view the list of users:.

```
show mgt-config users
users {
admin {
phash fnRL/G5lXVMug;
permissions {
role-based {
superuser yes;
}
}
}
admin_user_2 {
permissions {
role-based {
superreader yes;
}
}
authentication-profile RADIUS;
}
}
```

STEP 7 | Configure the Firewalls to Access the PAN-DB Private Cloud.

Configure the Firewalls to Access the PAN-DB Private Cloud

When using the PAN-DB public cloud, each firewall accesses the PAN-DB servers in the AWS cloud to download the list of eligible servers to which it can connect for URL lookups. With the PAN-DB private cloud, you must configure the firewalls with a (static) list of your PAN-DB private cloud servers that will be used for URL lookups. The list can contain up to 20 entries; IPv4 addresses, IPv6 addresses, and FQDNs are supported. Each entry on the list— IP address or FQDN—must be assigned to the management port and/or eth1 of the PAN-DB server.

STEP 1 | Pick one of the following options based on the PAN-OS version on the firewall.

- For firewalls running PAN-OS 7.0, access the PAN-OS CLI or the web interface on the firewall.

  Use the following CLI command to configure access to the private cloud:

```
set deviceconfig setting pan-url-db cloud-static-list <IP addresses>
 enable
```

  Or, in the web interface for each firewall, select **Device** > **Setup** > **Content-ID**, edit the URL Filtering section and enter the **PAN-DB Server** IP address(es) or FQDN(s). The list must be comma separated.

© 2018 Palo Alto Networks, Inc.

- For firewalls running PAN-OS 5.0, 6.0, or 6.1, use the following CLI command to configure access to the private cloud:

```
debug device-server pan-url-db cloud-static-list-enable <IP addresses>
 enable
```



To delete the entries for the private PAN-DB servers, and allow the firewalls to connect to the PAN-DB public cloud, use the command:

```
set deviceconfig setting pan-url-db cloud-static-list <IP addresses>
 disable
```

When you delete the list of private PAN-DB servers, a re-election process is triggered on the firewall. The firewall first checks for the list of PAN-DB private cloud servers and when it cannot find one, the firewall accesses the PAN-DB servers in the AWS cloud to download the list of eligible servers to which it can connect.

STEP 2 | **Commit** your changes.

STEP 3 | To verify that the change is effective, use the following CLI command on the firewall:

```
show url-cloud-status
Cloud status:            Up
URL database version:    20150417-220
```


© 2018 Palo Alto Networks, Inc.

URL Filtering Use Cases

The following use cases show how to use App-ID to control a specific set of web-based applications and how to use URL categories as match criteria in a policy. When working with App-ID, it is important to understand that each App-ID signature may have dependencies that are required to fully control an application. For example, with Facebook applications, the App#ID facebook#base is required to access the Facebook website and to control other Facebook applications. For example, to configure the firewall to control Facebook email, you would have to allow the App-IDs facebook-base and facebook-mail. As another example, if you search Applipedia (the App-ID database) for LinkedIn, you will see that in order to control LinkedIn mail, you need to apply the same action to both App-IDs: linkedin-base and linkedin-mail. To determine application dependencies for App-ID signatures, visit Applipedia, search for the given application, and then click the application for details.

- Use Case: Control Web Access on page 691
- Use Case: Use URL Categories for Policy Matching on page 694

These use cases rely on User-ID on page 445 to implement policies based on users and groups and a Decryption on page 589 to identify and control websites that are encrypted using SSL/TLS.

Use Case: Control Web Access

When using URL filtering to control user website access, there may be instances where granular control is required for a given website. In this use case, a URL filtering profile is applied to the security policy that allows web access for your users and the *social-networking* URL category is set to block, but the allow list in the URL profile is configured to allow the social networking site Facebook. To further control Facebook, the company policy also states that only marketing has full access to Facebook and all other users within the company can only read Facebook posts and cannot use any other Facebook applications, such as email, posting, chat, and file sharing. To accomplish this requirement, App-ID must be used to provide granular control over Facebook.

The first Security policy rule will allow marketing to access the Facebook website as well as all Facebook applications. Because this allow rule will also allow access to the Internet, threat prevention profiles are applied to the rule, so traffic that matches the policy will be scanned for threats. This is important because the allow rule is terminal and will not continue to check other rules if there is a traffic match.

STEP 1 | Confirm that URL filtering is licensed.

1. Select **Device** > **Licenses** and confirm that a valid date appears for the URL filtering database that will used. This will either be PAN-DB or BrightCloud.
2. If a valid license is not installed, see Enable PAN-DB URL Filtering.

STEP 2 | Confirm that User-ID is working. User-ID is required to create policies based on users and groups.

1. To check Group Mapping from the CLI, enter the following command:

   show user group-mapping statistics
2. To check User Mapping from the CLI, enter the following command:

   show user ip-user-mapping-mp all
3. If statistics do not appear and/or IP address to user mapping information is not displayed, see User-ID.

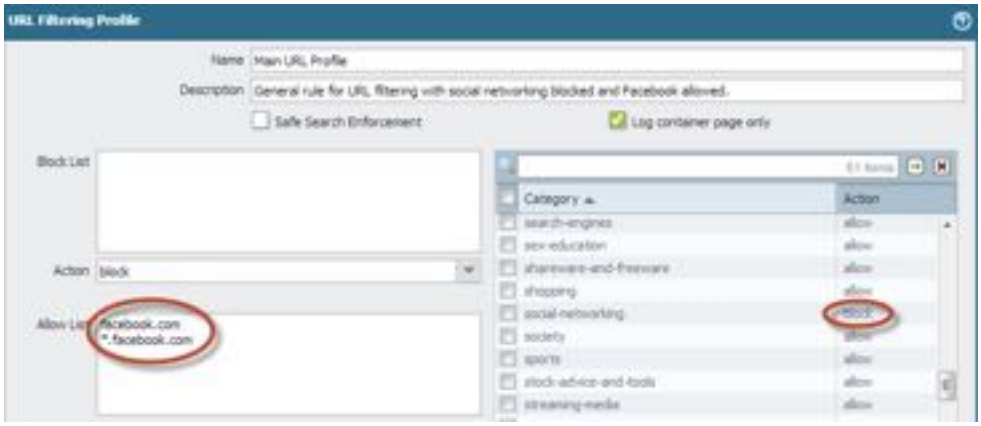
STEP 3 | Set up a URL filtering profile by cloning the default profile.

© 2018 Palo Alto Networks, Inc.

1. Select **Objects** > **Security Profiles** > **URL Filtering** and select the **default** profile.
2. Click the **Clone** icon. A new profile should appear named **default-1**.
3. Select the new profile and rename it.

STEP 4 | Configure the URL filtering profile to block social-networking and allow Facebook.

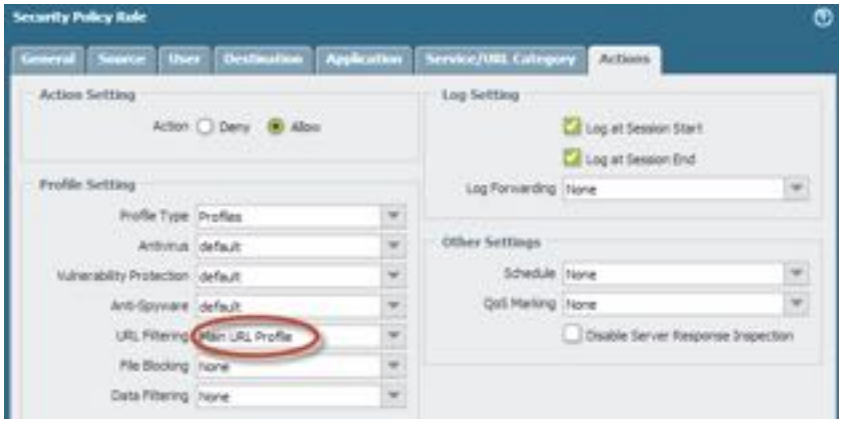
1. Modify the new URL filtering profile and in the **Category** list scroll to **social-networking** and in the **Action** column click on **allow** and change the action to **block**.
2. In the **Allow List**, enter **`facebook.com`**, press enter to start a new line and then type **`*.facebook.com`**. Both of these formats are required, so all URL variants a user may use will be identified, such as facebook.com, www.facebook.com, and https://facebook.com.



3. Click **OK** to save the profile.

STEP 5 | Apply the new URL filtering profile to the security policy rule that allows web access from the user network to the Internet.

1. Select **Policies** > **Security** and click on the policy rule that allows web access.
2. On the **Actions** tab, select the URL profile you just created from the **URL Filtering** drop-down.



3. Click **OK** to save.

STEP 6 | Create the security policy rule that will allow marketing access the Facebook website and all Facebook applications.

This rule must precede other rules because:

- It is a specific rule. More specific rules must precede other rules.
- Allow rule will terminate when a traffic match occurs.

1. Select **Policies** > **Security** and click **Add**.

© 2018 Palo Alto Networks, Inc.

2. Enter a **Name** and optionally a **Description** and **Tag**(s).
3. On the **Source** tab add the zone where the users are connected.
4. On the **User** tab in the **Source User** section click **Add**.
5. Select the directory group that contains your *marketing* users.
6. On the **Destination** tab, select the zone that is connected to the Internet.
7. On the **Applications** tab, click **Add** and add the *facebook* App-ID signature.
8. On the **Actions** tab, add the default profiles for **Antivirus**, **Vulnerability Protection**, and **Anti-Spyware**.
9. Click **OK** to save the security profile.

   The *facebook* App-ID signature used in this policy rule encompasses all Facebook applications, such as facebook-base, facebook-chat, and facebook-mail, so this is the only App-ID signature required in this rule.

   With this rule in place, when a marketing employee attempts to access the Facebook website or any Facebook application, the rule matches based on the user being part of the marketing group. For traffic from any user outside of marketing, the rule will be skipped because there would not be a traffic match and rule processing would continue.

STEP 7 | Configure the security policy to block all other users from using any Facebook applications other than simple web browsing. The easiest way to do this is to clone the marketing allow policy and then modify it.

1. From **Policies** > **Security** click the marketing Facebook allow policy you created earlier to highlight it and then click the **Clone** icon.
2. Enter a **Name** and optionally enter a **Description** and **Tag**('s).
3. On the **User** tab highlight the marketing group and delete it and in the drop-down select **any**.
4. On the **Applications** tab, click the *facebook* App-ID signature and delete it.
5. Click **Add** and add the following App-ID signatures:
   - facebook-apps
   - facebook-chat
   - facebook-file-sharing
   - facebook-mail
   - facebook-posting
   - facebook-social-plugin
6. On the **Actions** tab in the **Action Setting** section, select **Deny**. The profile settings should already be correct because this rule was cloned.



7. Click **OK** to save the security profile.
8. Ensure that this new deny rule is listed after the marketing allow rule, to ensure that rule processing occurs in the correct order to allow marketing users and then to deny/limit all other users.

© 2018 Palo Alto Networks, Inc.

9. Click **Commit** to save the configuration.

   With these security policy rules in place, any user who is part of the marketing group will have full access to all Facebook applications and any user that is not part of the marketing group will only have read-only access to the Facebook website and will not be able to use Facebook applications such as post, chat, email, and file sharing.

Use Case: Use URL Categories for Policy Matching

You can also use URL categories as match criteria in the following policy types: Authentication, Decryption, Security, and QoS. In this use case, Decryption policy rules match on URL categories to control which web categories to decrypt or not decrypt. The first rule is a no-decrypt rule instructing the firewall not to decrypt outbound user traffic to financial-services or health-and-medicine sites and the second rule instructs the firewall to decrypt all other traffic.

STEP 1 | Create the no-decrypt rule that will be listed first in the decryption policies list. This will prevent any website that is in the ***financial-services*** or ***health-and-medicine*** URL categories from being decrypted.

1. Select **Policies** > **Decryption** and click **Add**.
2. Enter a **Name** and optionally enter a **Description** and **Tag**(s).
3. On the **Source** tab, add the zone where the users are connected.
4. On the **Destination** tab, enter the zone that is connected to the Internet.
5. On the **URL Category** tab, click **Add** and select the *financial-services* and *health-and-medicine* URL categories.
6. On the **Options** tab, set the action to **No Decrypt**.
7. (Optional) Although the firewall does not decrypt and inspect the traffic for the session, you can attach a **Decryption profile** if you want to enforce the server certificates used during the session. The decryption profile allows you to configure the firewall to terminate the SSL connection either when the server certificates are expired or when the server certificates are issues by an untrusted issuer.



8. Click **OK** to save the policy rule.

STEP 2 | Create the decryption policy rule that will decrypt all other traffic.

1. Select the no-decrypt policy you created previously and then click **Clone**.
2. Enter a **Name** and optionally enter a **Description** and **Tag**(s).
3. On the **URL Category** tab, select *financial-services* and *health-and-medicine* and then click the **Delete** icon.
4. On the **Options** tab, set the action to **Decrypt** and the **Type** to **SSL Forward Proxy**.
5. (Optional) Attach a **Decryption profile** to specify the server certificate verification, unsupported mode checks and failure checks for the SSL traffic. See Configure SSL Forward Proxy for more details.

© 2018 Palo Alto Networks, Inc.



6. Ensure that this new decryption rule is listed after the no-decrypt rule to ensure that rule processing occurs in the correct order, so websites in the *financial-services* and *health-and-medicine* are not decrypted
7. Click **OK** to save the policy rule.

STEP 3 | (BrightCloud only) Enable cloud lookups for dynamically categorizing a URL when the category is not available on the local database on the firewall.

1. Access the CLI on the firewall.
2. Enter the following commands to enable Dynamic URL Filtering:
    1. **configure**
    2. **set deviceconfig setting url dynamic-url yes**
    3. **commit**

STEP 4 | Save the configuration.

Click **Commit**.

With these two decrypt policies in place, any traffic destined for the *financial-services* or *health-and-medicine* URL categories will not be decrypted. All other traffic will be decrypted.

Now that you have a basic understanding of the powerful features of URL filtering, App-ID, and User-ID, you can apply similar policies to your firewall to control any application in the Palo Alto Networks App-ID signature database and control any website contained in the URL filtering database.

For help in troubleshooting URL filtering issues, see Troubleshoot URL Filtering.

© 2018 Palo Alto Networks, Inc.

Troubleshoot URL Filtering

The following topics provide troubleshooting guidelines for diagnosing and resolving common URL filtering problems.

- Problems Activating PAN-DB on page 696
- PAN-DB Cloud Connectivity Issues on page 696
- URLs Classified as Not-Resolved on page 698
- Incorrect Categorization on page 698
- URL Database Out of Date on page 699

Problems Activating PAN-DB

Use the following workflow to troubleshoot PAN-DB activation issues.

STEP 1 | Access the PAN-OS CLI.

STEP 2 | Verify whether PAN-DB has been activated by running the following command:

```
show system setting url-database
```

If the response is `paloaltonetworks`, PAN-DB is the active vendor.

STEP 3 | Verify that the firewall has a valid PAN-DB license by running the following command:

```
request license info
```

You should see the license entry `Feature: PAN_DB URL Filtering`. If the license is not installed, you will need to obtain and install a license. See Configure URL Filtering.

STEP 4 | After installing the license, download a new PAN-DB seed database by running the following command:

```
request url-filtering download paloaltonetworks region <region>
```

STEP 5 | Check the download status by running the following command:

```
request url-filtering download status vendor paloaltonetworks
```

- If the message is different from `PAN-DB download: Finished successfully`, stop here; there may be a problem connecting to the cloud. Attempt to solve the connectivity issue by performing basic network troubleshooting between the firewall and the Internet. For more information, see PAN-DB Cloud Connectivity Issues.
- If the message is `PAN-DB download: Finished successfully`, the firewall successfully downloaded the URL seed database. Try to enable PAN-DB again by running the following command:

```
admin@PA-200> set system setting url-database paloaltonetworks
```

- If the problems persists, contact Palo Alto Networks Customer Support.

PAN-DB Cloud Connectivity Issues

To check connectivity between the firewall and the PAN-DB cloud:

```
show url-cloud status
```


© 2018 Palo Alto Networks, Inc.

If the cloud is accessible, the expected response is similar to the following:

```
show url-cloud status
PAN-DB URL Filtering
License :                           valid
Current cloud server :               s0000.urlcloud.paloaltonetworks.com
Cloud connection :                  connected
URL database version - device :     2013.11.18.000
URL database version - cloud :      2013.11.18.000  ( last update time
 2013/11/19
 13:20:51 )
URL database status :               good
URL protocol version - device :     pan/0.0.2
URL protocol version - cloud :      pan/0.0.2
Protocol compatibility status :     compatible
```

If the cloud is not accessible, the expected response is similar to the following:

```
show url-cloud status
PAN-DB URL Filtering
License :                           valid
Cloud connection :                  not connected
URL database version - device :     2013.11.18.000
URL database version - cloud :      2013.11.18.000  ( last update time
 2013/11/19
 13:20:51 )
URL database status :               good
URL protocol version - device :     pan/0.0.2
URL protocol version - cloud :      pan/0.0.2
Protocol compatibility status :     compatible
```

Use the following checklist to identify and resolve connectivity issues:

- Does the PAN-DB URL Filtering license field shows as invalid? Obtain and install a valid PAN-DB license.
- Does the URL database status show as out of date? Download a new seed database by running the following command:

```
request url-filtering download paloaltonetworks region <region>
```

- Does the URL protocol version show as not compatible? Upgrade PAN-OS to the latest version.
- Can you ping the PAN-DB cloud server from the firewall? Run the following command to check:

```
ping source <ip-address> host s0000.urlcloud.paloaltonetworks.com <
```

  For example, if your management interface IP address is 10.1.1.5, run the following command:

```
ping source 10.1.1.5 host s0000.urlcloud.paloaltonetworks.com
```

- Is the firewall in an HA configuration? Verify that the HA state of the firewalls is in the active, active-primary, or active-secondary state. Access to the PAN-DB cloud will be blocked if the firewall is in a different state. Run the following command on each firewall in the pair to see the state:

```
show high-availability state
```

If you still have problems with connectivity between the firewall and the PAN-DB cloud, contact Palo Alto Networks support.

© 2018 Palo Alto Networks, Inc.

URLs Classified as Not-Resolved

Use the following workflow to troubleshoot why some or all of the URLs being identified by PAN-DB are classified as Not-resolved:

STEP 1 | Check the PAN-DB cloud connection by running the following command:

```
show url-cloud status
```

The Cloud connection: field should show `connected`. If you see anything other than `connected`, any URL that do not exist in the management plane cache will be categorized as `not-resolved`. To resolve this issue, see PAN-DB Cloud Connectivity Issues.

STEP 2 | If the cloud connection status shows **`connected`**, check the current utilization of the firewall. If firewall utilization is spiking, URL requests may be dropped (may not reach the management plane), and will be categorized as **`not-resolved`**.

To view system resources, run the following command and view the `%CPU` and `%MEM` columns:

```
show system resources
```

You can also view system resources on the System Resources widget on the **Dashboard** in the web interface.

STEP 3 | If the problem persist, contact Palo Alto Networks support.

Incorrect Categorization

Sometimes you may come across a URL that you believe is categorized incorrectly. Use the following workflow to determine the URL categorization for a site and request a category change, if appropriate.

STEP 1 | Verify the category in the dataplane by running the following command:

```
show running url <URL>
```

For example, to view the category for the Palo Alto Networks website, run the following command:

```
show running url paloaltonetworks.com
```

If the URL stored in the dataplane cache has the correct category (computer-and-internet-info in this example), then the categorization is correct and no further action is required. If the category is not correct, continue to the next step.

STEP 2 | Verify if the category in the management plane by running the command:

```
test url-info-host <URL>
```

For example:

```
test url-info-host paloaltonetworks.com
```

If the URL stored in the management plane cache has the correct category, remove the URL from the dataplane cache by running the following command:

© 2018 Palo Alto Networks, Inc.

```
clear url-cache url <URL>
```

The next time the firewall requests the category for this URL, the request will be forwarded to the management plane. This will resolve the issue and no further action is required. If this does not solve the issue, go to the next step to check the URL category on the cloud systems.

STEP 3 | Verify the category in the cloud by running the following command:

```
test url-info-cloud <URL>
```

STEP 4 | If the URL stored in the cloud has the correct category, remove the URL from the dataplane and the management plane caches.

Run the following command to delete a URL from the dataplane cache:

```
clear url-cache url <URL>
```

Run the following command to delete a URL from the management plane cache:

```
delete url-database url <URL>
```

The next time the firewall queries for the category of the given URL, the request will be forwarded to the management plane and then to the cloud. This should resolve the category lookup issue. If problems persist, see the next step to submit a categorization change request.

STEP 5 | To submit a change request from the web interface, go to the URL log and select the log entry for the URL you would like to have changed.

STEP 6 | Click the **Request Categorization** change link and follow instructions. You can also request a category change from the Palo Alto Networks Test A Site website by searching for the URL and then clicking the **Request Change** icon. To view a list of all available categories with descriptions of each category, refer to https://urlfiltering.paloaltonetworks.com/CategoryList.aspx.

If your change request is approved, you will receive an email notification. You then have two options to ensure that the URL category is updated on the firewall:

- Wait until the URL in the cache expires and the next time the URL is accessed by a user, the new categorization update will be put in the cache.
- Run the following command to force an update in the cache:

```
request url-filtering update url <URL>
```

URL Database Out of Date

If you have observed through the syslog or the CLI that PAN-DB is out-of-date, it means that the connection from the firewall to the PAN-DB cloud is blocked. This usually occurs when the URL database on the firewall is too old (version difference is more than three months) and the cloud cannot update the firewall automatically. In order to resolve this issue, you must re-download an initial seed database (this operation is not blocked). This will result in an automatic re-activation of PAN-DB.

© 2018 Palo Alto Networks, Inc.

To manually update the database, perform one of the following steps:

- From the web interface, select **Device** > **Licenses** and in the **PAN-DB URL Filtering** section click the **Re-Download** link.
- From the CLI, run the following command:

```
request url-filtering download paloaltonetworks region <region_name>
```



Re-downloading the seed database causes the URL cache in the management plane and dataplane to be purged. The management plane cache will then be re-populated with the contents of the new seed database.

Quality of Service

Quality of Service (QoS) is a set of technologies that work on a network to guarantee its ability to dependably run high-priority applications and traffic under limited network capacity. QoS technologies accomplish this by providing differentiated handling and capacity allocation to specific flows in network traffic. This enables the network administrator to assign the order in which traffic is handled, and the amount of bandwidth afforded to traffic.

Palo Alto Networks Application Quality of Service (QoS) provides basic QoS applied to networks and extends it to provide QoS to applications and users.

Use the following topics to learn about and configure Palo Alto Networks application-based QoS:

- QoS Overview
- QoS Concepts
- Configure QoS
- Configure QoS for a Virtual System
- Enforce QoS Based on DSCP Classification
- QoS Use Cases

Use the Palo Alto Networks product comparison tool to view the QoS features supported on your firewall model. Select two or more product models and click **Compare Now** to view QoS feature support for each model (for example, you can check if your firewall model supports QoS on subinterfaces and if so, the maximum number of subinterfaces on which QoS can be enabled).

QoS on Aggregate Ethernet (AE) interfaces is supported on PA-7000 Series, PA-5000 Series, PA-3000 Series, and PA-5200 Series firewalls running PAN-OS 7.0 or later release versions.

© 2018 Palo Alto Networks, Inc.

QoS Overview

Use QoS to prioritize and adjust quality aspects of network traffic. You can assign the order in which packets are handled and allot bandwidth, ensuring preferred treatment and optimal levels of performance are afforded to selected traffic, applications, and users.

Service quality measurements subject to a QoS implementation are bandwidth (maximum rate of transfer), throughput (actual rate of transfer), latency (delay), and jitter (variance in latency). The capability to shape and control these service quality measurements makes QoS of particular importance to high-bandwidth, real-time traffic such as voice over IP (VoIP), video conferencing, and video-on-demand that has a high sensitivity to latency and jitter. Additionally, use QoS to achieve outcomes such as the following:

- Prioritize network and application traffic, guaranteeing high priority to important traffic or limiting non-essential traffic.
- Achieve equal bandwidth sharing among different subnets, classes, or users in a network.
- Allocate bandwidth externally or internally or both, applying QoS to both upload and download traffic or to only upload or download traffic.
- Ensure low latency for customer and revenue-generating traffic in an enterprise environment.
- Perform traffic profiling of applications to ensure bandwidth usage.

QoS implementation on a Palo Alto Networks firewall begins with three primary configuration components that support a full QoS solution: a QoS Profile, a QoS Policy, and setting up the QoS Egress Interface. Each of these options in the QoS configuration task facilitate a broader process that optimizes and prioritizes the traffic flow and allocates and ensures bandwidth according to configurable parameters.

The figure QoS Traffic Flow shows traffic as it flows from the source, is shaped by the firewall with QoS enabled, and is ultimately prioritized and delivered to its destination.

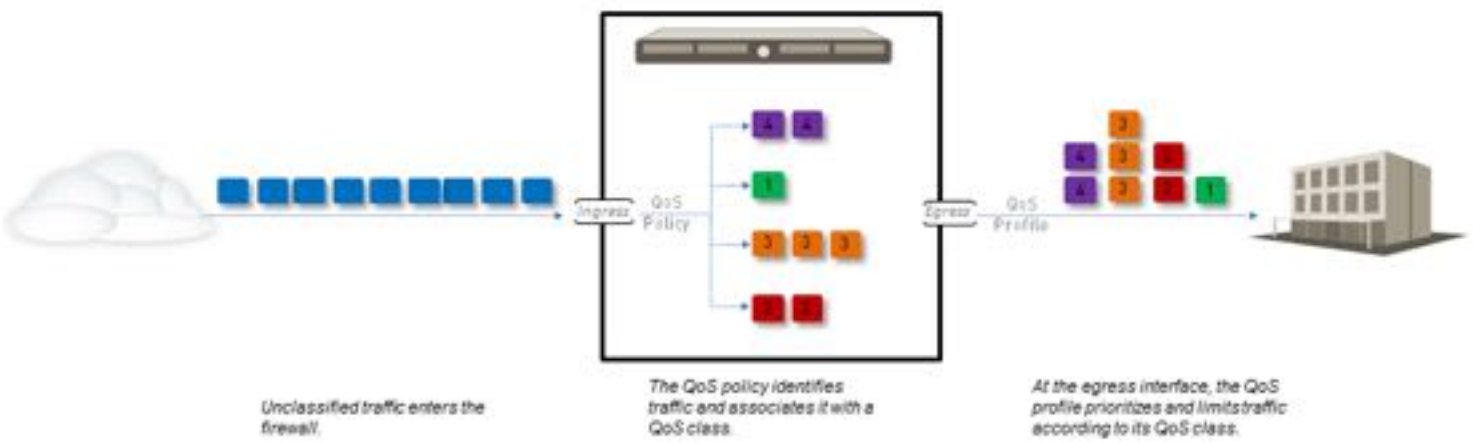


Figure 8: QoS Traffic Flow

The QoS configuration options allow you to control the traffic flow and define it at different points in the flow. The figure QoS Traffic Flow indicates where the configurable options define the traffic flow. A QoS policy rule allows you to define traffic you want to receive QoS treatment and assign that traffic a QoS class. The matching traffic is then shaped based on the QoS profile class settings as it exits the physical interface.

Each of the QoS configuration components influence each other and the QoS configuration options can be used to create a full and granular QoS implementation or can be used sparingly with minimal administrator action.

© 2018 Palo Alto Networks, Inc.

Each firewall model supports a maximum number of ports that can be configured with QoS. Refer to the spec sheet for your firewall model or use the product comparison tool to view QoS feature support for two or more firewalls on a single page.

© 2018 Palo Alto Networks, Inc.

QoS Concepts

Use the following topics to learn about the different components and mechanisms of a QoS configuration on a Palo Alto Networks firewall:

- QoS for Applications and Users on page 705
- QoS Policy on page 705
- QoS Profile on page 705
- QoS Classes on page 706
- QoS Priority Queuing on page 706
- QoS Bandwidth Management on page 706
- QoS Egress Interface on page 707
- QoS for Clear Text and Tunneled Traffic on page 708

QoS for Applications and Users

A Palo Alto Networks firewall provides basic QoS, controlling traffic leaving the firewall according to network or subnet, and extends the power of QoS to also classify and shape traffic according to application and user. The Palo Alto Networks firewall provides this capability by integrating the features App-ID on page 509 and User-ID on page 445 with the QoS configuration. App-ID and User-ID entries that exist to identify specific applications and users in your network are available in the QoS configuration so that you can easily specify applications and users for which you want to manage and/or guarantee bandwidth.

QoS Policy

Use a QoS policy rule to define traffic to receive QoS treatment (either preferential treatment or bandwidth-limiting) and assigns such traffic a QoS class of service.

Define a QoS policy rule to match to traffic based on:

- Applications and application groups.
- Source zones, source addresses, and source users.
- Destination zones and destination addresses.
- Services and service groups limited to specific TCP and/or UDP port numbers.
- URL categories, including custom URL categories.
- Differentiated Services Code Point (DSCP) and Type of Service (ToS) values, which are used to indicate the level of service requested for traffic, such as high priority or best effort delivery.

Set up multiple QoS policy rules (**Policies** > **QoS**) to associate different types of traffic with different QoS Classes on page 706 of service.

QoS Profile

Use a QoS profile rule to define values of up to eight QoS Classes contained within that single profile rule.

With a QoS profile rule, you can define QoS Priority Queuing and QoS Bandwidth Management for QoS classes. Each QoS profile rule allows you to configure individual bandwidth and priority settings for up eight QoS classes, as well as the total bandwidth alloted for the eight classes combined. Attach the QoS profile rule (or multiple QoS profile rules) to a physical interface to apply the defined priority and bandwidth settings to the traffic exiting that interface.

A default QoS profile rule is available on the firewall. The default profile rule and the classes defined in the profile do not have predefined maximum or guaranteed bandwidth limits.

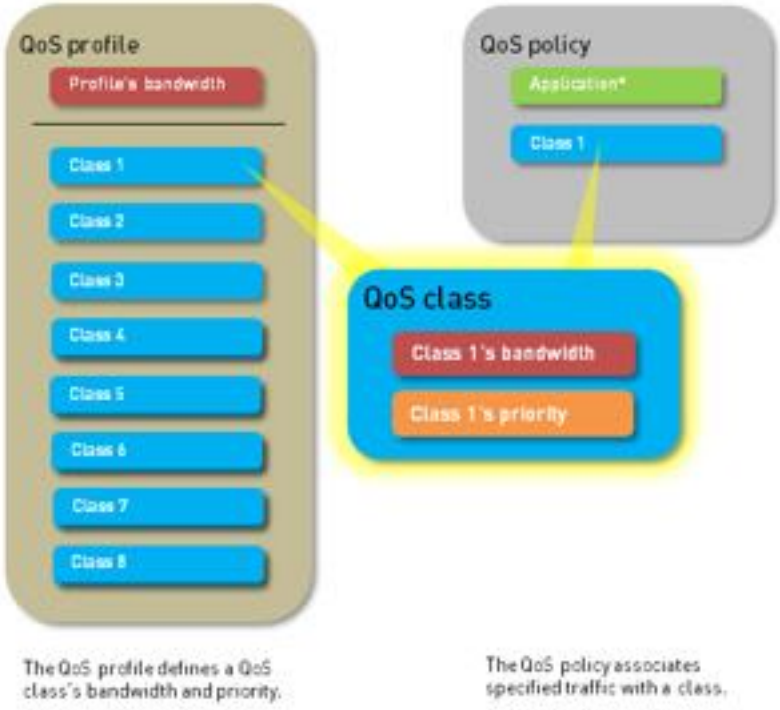
© 2018 Palo Alto Networks, Inc.

To define priority and bandwidth settings for QoS classes, see Step Add a QoS profile rule.

QoS Classes

A QoS class determines the priority and bandwidth for traffic matching a QoS Policy rule. You can use a QoS Profile rule to define QoS classes. There are up to eight definable QoS classes in a single QoS profile. Unless otherwise configured, traffic that does not match a QoS class is assigned a class of 4.

QoS Priority Queuing and QoS Bandwidth Management, the fundamental mechanisms of a QoS configuration, are configured within the QoS class definition (see Step Add a QoS profile rule.). For each QoS class, you can set a priority (real-time, high, medium, and low) and the maximum and guaranteed bandwidth for matching traffic. QoS priority queuing and bandwidth management determine the order of traffic and how traffic is handled upon entering or leaving a network.



QoS Priority Queuing

One of four priorities can be enforced for a QoS class: real-time, high, medium, and low. Traffic matching a QoS policy rule is assigned the QoS class associated with that rule, and the firewall treats the matching traffic based on the QoS class priority. Packets in the outgoing traffic flow are queued based on their priority until the network is ready to process the packets. Priority queuing allows you to ensure that important traffic, applications, and users take precedence. Real-time priority is typically used for applications that are particularly sensitive to latency, such as voice and video applications.

QoS Bandwidth Management

QoS bandwidth management allows you to control traffic flows on a network so that traffic does not exceed network capacity (resulting in network congestion) and 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 a 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

© 2018 Palo Alto Networks, Inc.

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 for traffic matching that QoS class (QoS classes are assigned to traffic matching QoS Policy rules) and the overall bandwidth limit for the profile can be applied to all clear text traffic, specific clear text traffic originating from source interfaces and source subnets, all 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.

The following fields support QoS bandwidth settings:

- **Egress Guaranteed**—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 traffic. Depending on your QoS configuration, you can guarantee bandwidth for a single QoS class, for all or some clear text traffic, and for all or some tunneled traffic.

  Example:

  Class 1 traffic has 5 Gbps of egress guaranteed bandwidth, which means that 5 Gbps 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, 5 Gbps of bandwidth is absolutely available for class 1 traffic. During these periods of congestion, any Class 1 traffic that exceeds 5 Gbps is best effort.
- **Egress Max**—The overall bandwidth allocation for matching traffic. The firewall drops traffic that exceeds the egress max limit that you set. Depending on your QoS configuration, you can set a maximum bandwidth limit for a QoS class, for all or some clear text traffic, for all or some tunneled traffic, and for all traffic exiting the QoS interface.



The cumulative guaranteed bandwidth for the QoS profile rules attached to the interface must not exceed the total bandwidth allocated to the interface.

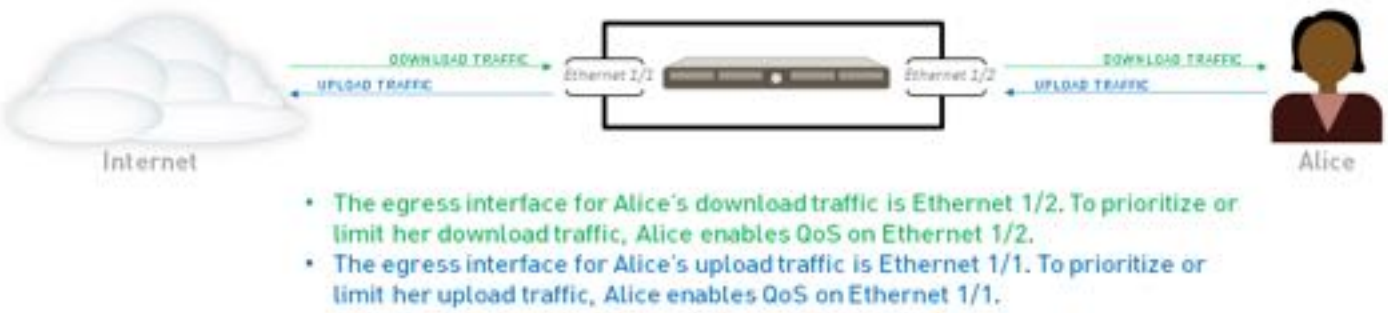
To define bandwidth settings for QoS classes, see Step Add a QoS profile rule. To then apply those bandwidth settings to clear text and tunneled traffic, and to set the overall bandwidth limit for a QoS interface, see Step Enable QoS on a physical interface.

QoS Egress Interface

Enabling a QoS profile rule on the egress interface of the traffic identified for QoS treatment completes a QoS configuration. The ingress interface for QoS traffic is the interface on which the traffic enters the firewall. The egress interface for QoS traffic is the interface that traffic leaves the firewall from. QoS is always enabled and enforced on the egress interface for a traffic flow. The egress interface in a QoS configuration can either be the external- or internal-facing interface of the firewall, depending on the flow of the traffic receiving QoS treatment.

For example, in an enterprise network, if you are limiting employees' download traffic from a specific website, the egress interface in the QoS configuration is the firewall's internal interface, as the traffic flow is from the Internet, through the firewall, and to your company network. Alternatively, when limiting employees' upload traffic to the same website, the egress interface in the QoS configuration is the firewall's external interface, as the traffic you are limiting flows from your company network, through the firewall, and then to the Internet.

© 2018 Palo Alto Networks, Inc.



See Step Identify the egress interface for applications that you identified as needing QoS treatment. to learn how to Identify the egress interface for applications that you want to receive QoS treatment.

QoS for Clear Text and Tunneled Traffic

At the minimum, enabling a QoS interfaces requires you to select a default QoS profile rule that defines bandwidth and priority settings for clear text traffic egressing the interface. However, when setting up or modifying a QoS interface, you can apply granular QoS settings to outgoing clear text traffic and tunneled traffic. QoS preferential treatment and bandwidth limiting can be enforced for tunneled traffic, for individual tunnel interfaces, and/or for clear text traffic originating from different source interfaces and source subnets. On Palo Alto Networks firewalls, *tunneled traffic* refers to tunnel interface traffic, specifically IPSec traffic in tunnel mode.

© 2018 Palo Alto Networks, Inc.

Configure QoS

Follow these steps to configure Quality of Service (QoS), which includes creating a QoS profile, creating a QoS policy, and enabling QoS on an interface.

STEP 1 | Identify the traffic you want to manage with QoS.

This example shows how to use QoS to limit web browsing.

Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.

Click any application name to display detailed application information.

STEP 2 | Identify the egress interface for applications that you want to receive QoS treatment.

The egress interface for traffic depends on the traffic flow. If you are shaping incoming traffic, the egress interface is the internal-facing interface. If you are shaping outgoing traffic, the egress interface is the external-facing interface.

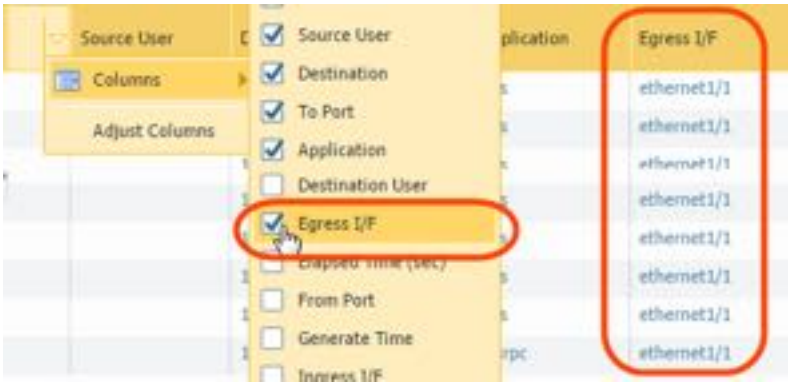
Select **Monitor** > **Logs** > **Traffic** to view the Traffic logs.

To filter and only show logs for a specific application:

- If an entry is displayed for the application, click the underlined link in the Application column then click the Submit icon.
- If an entry is not displayed for the application, click the Add Log icon and search for the application.

The **Egress I/F** in the traffic logs displays each application's egress interface. To display the **Egress I/F** column if it is not displayed by default:

- Click any column header to add a column to the log:



- Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface listed in the Destination section:

© 2018 Palo Alto Networks, Inc.



STEP 3 | Add a QoS policy rule.

A QoS policy rule defines the traffic to receive QoS treatment. The firewall assigns a QoS class of service to the traffic matched to the policy rule.

1. Select **Policies** > **QoS** and **Add** a new policy rule.
2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**.
3. Specify traffic to receive QoS treatment based on **Source**, **Destination**, **Application**, **Service/URL Category**, and **DSCP/ToS** values (the **DSCP/ToS** settings allow you to Enforce QoS Based on DSCP Classification).

   For example, select the **Application**, click **Add**, and select **web-browsing** to apply QoS to web browsing traffic.
4. (Optional) Continue to define additional parameters. For example, select **Source** and **Add** a **Source User** to provide QoS for a specific user's web traffic.
5. Select **Other Settings** and assign a **QoS Class** to traffic matching the policy rule. For example, assign Class 2 to the user1's web traffic.
6. Click **OK**.

STEP 4 | Add a QoS profile rule.

A QoS profile rule allows you to define the eight classes of service that traffic can receive, including priority, and enables QoS Bandwidth Management.

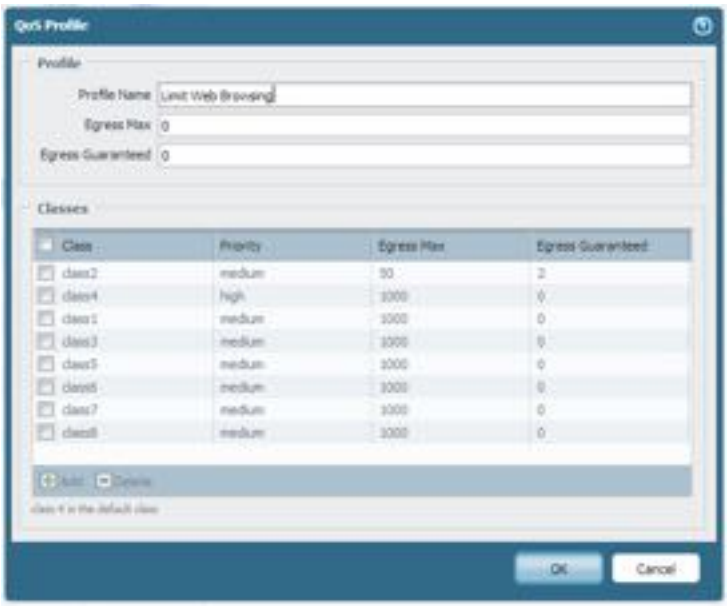
You can edit any existing QoS profile, including the default, by clicking the QoS profile name.

1. Select **Network** > **Network Profiles** > **QoS Profile** and **Add** a new profile.
2. Enter a descriptive **Profile Name**.
3. Set the overall bandwidth limits for the QoS profile rule:
   - Enter an **Egress Max** value to set the overall bandwidth allocation for the QoS profile rule.
   - Enter an **Egress Guaranteed** value to set the guaranteed bandwidth for the QoS Profile.

   Any traffic that exceeds the Egress Guaranteed value is best effort and not guaranteed. Bandwidth that is guaranteed but is unused continues to remain available for all traffic.
4. In the Classes section, specify how to treat up to eight individual QoS classes:
   1. **Add** a class to the QoS Profile.
   2. Select the **Priority** for the class: real-time, high, medium, or low.
   3. Enter the **Egress Max** and **Egress Guaranteed** bandwidth for traffic assigned to each QoS class.
5. Click **OK**.

In the following example, the QoS profile rule Limit Web Browsing limits Class 2 traffic to a maximum bandwidth of 50Mbps and a guaranteed bandwidth of 2Mbps.

© 2018 Palo Alto Networks, Inc.



STEP 5 | Enable QoS on a physical interface.

Part of this step includes the option to select clear text and tunneled traffic for unique QoS treatment.

Check if the firewall model you're using supports enabling QoS on a subinterface by reviewing a summary of the Product Specifications.

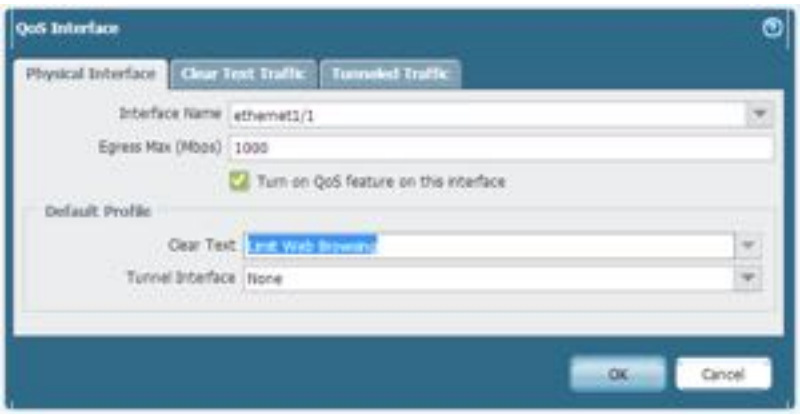
1. Select **Network** > **QoS** and **Add** a QoS interface.
2. Select **Physical Interface** and choose the **Interface Name** of the interface on which to enable QoS.

   In the example, Ethernet 1/1 is the egress interface for web-browsing traffic (see Step 2).
3. Set the **Egress Max** bandwidth for all traffic exiting this interface.

   It is a best practice to always define the Egress Max value for a QoS interface. Ensure that the cumulative guaranteed bandwidth for the QoS profile rules attached to the interface does not exceed the total bandwidth allocated to the interface.
4. Select **Turn on QoS feature on this interface**.
5. In the Default Profile section, select a QoS profile rule to apply to all **Clear Text** traffic exiting the physical interface.
6. (Optional) Select a default QoS profile rule to apply to all tunneled traffic exiting the interface.

For example, enable QoS on ethernet 1/1 and apply the bandwidth and priority settings you defined for the QoS profile rule Limit Web Browsing (Step 4) to be used as the default settings for clear text egress traffic.

© 2018 Palo Alto Networks, Inc.



1. (Optional) Continue to define more granular settings to provide QoS for Clear Text and Tunneled Traffic. Settings configured on the **Clear Text Traffic** tab and the **Tunneled Traffic** tab automatically override the default profile settings for clear text and tunneled traffic on the Physical Interface tab.
   - Select **Clear Text Traffic** and:
     - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
     - Click **Add** and apply a QoS profile rule to enforce clear text traffic based on source interface and source subnet.
   - Select **Tunneled Traffic** and:
     - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
     - Click **Add** and attach a QoS profile rule to a single tunnel interface.
2. Click **OK**.

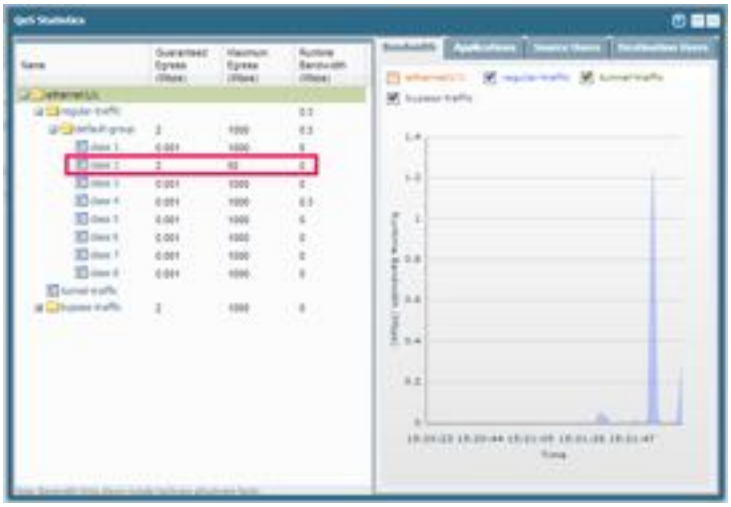
STEP 6 | Commit your changes.

Click **Commit**.

STEP 7 | Verify a QoS configuration.

Select **Network** > **QoS** and then **Statistics** to view QoS bandwidth, active sessions of a selected QoS class, and active applications for the selected QoS class.

For example, see the statistics for ethernet 1/1 with QoS enabled:



Class 2 traffic limited to 2Mbps of guaranteed bandwidth and a maximum bandwidth of 50Mbps.

© 2018 Palo Alto Networks, Inc.

Continue to click the tabs to display further information regarding applications, source users, destination users, security rules and QoS rules.



Bandwidth limits shown on the QoS Statistics window include a hardware adjustment factor.

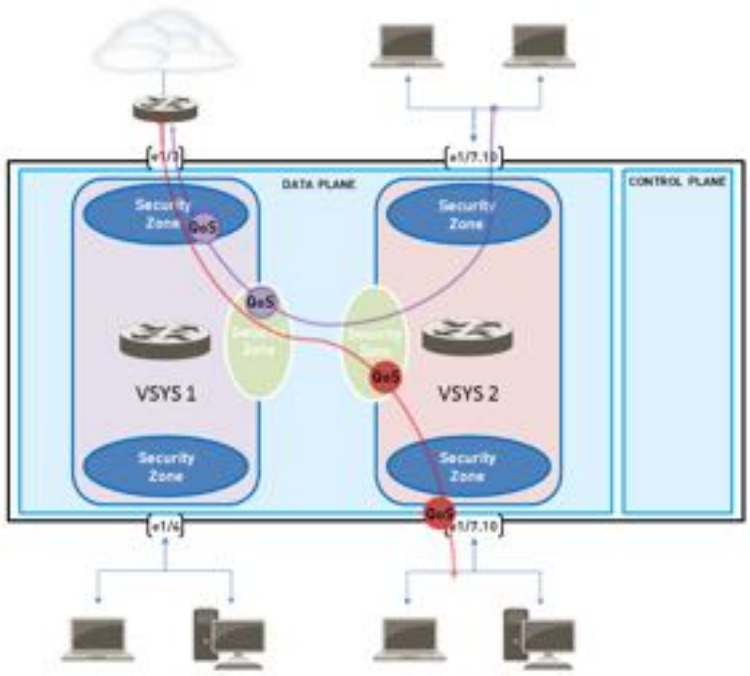
© 2018 Palo Alto Networks, Inc.

Configure QoS for a Virtual System

QoS can be configured for a single or several virtual systems configured on a Palo Alto Networks firewall. Because a virtual system is an independent firewall, QoS must be configured independently for a single virtual system.

Configuring QoS for a virtual system is similar to configuring QoS on a physical firewall, with the exception that configuring QoS for a virtual system requires specifying the source and destination of traffic. Because a virtual system exists without set physical boundaries and because traffic in a virtual environment spans more than one virtual system, specifying source and destination zones and interfaces for traffic is necessary to control and shape traffic for a single virtual system.

The example below shows two virtual systems configured on firewall. VSYS 1 (purple) and VSYS 2 (red) each have QoS configured to prioritize or limit two distinct traffic flows, indicated by their corresponding purple (VSYS 1) and red (VSYS 2) lines. The QoS nodes indicate the points at traffic is matched to a QoS policy and assigned a QoS class of service, and then later indicate the point at which traffic is shaped as it egresses the firewall.



Refer to Virtual Systems for information on virtual systems and how to configure them.

STEP 1 | Confirm that the appropriate interfaces, virtual routers, and security zones are associated with each virtual system.

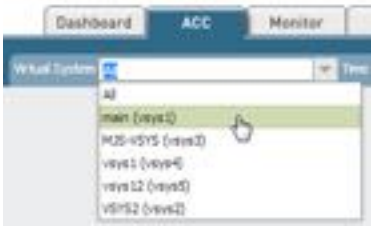
- To view configured interfaces, select **Network** > **Interface**.
- To view configured zones, select **Network** > **Zones**.
- To view information on defined virtual routers, select **Network** > **Virtual Routers**.

STEP 2 | Identify traffic to apply QoS to.

Select **ACC** to view the **Application Command Center** page. Use the settings and charts on the **ACC** page to view trends and traffic related to Applications, URL filtering, Threat Prevention, Data Filtering, and HIP Matches.

© 2018 Palo Alto Networks, Inc.

To view information for a specific virtual system, select the virtual system from the **Virtual System** drop-down:



Click any application name to display detailed application information.

STEP 3 | Identify the egress interface for applications that you identified as needing QoS treatment.

In a virtual system environment, QoS is applied to traffic on the traffic's egress point on the virtual system. Depending the configuration and QoS policy for a virtual system, the egress point of QoS traffic could be associated with a physical interface or could be a zone.

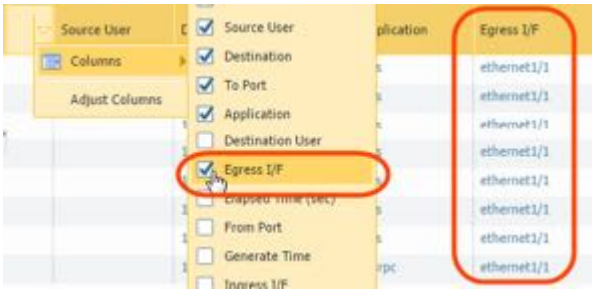
This example shows how to limit web-browsing traffic on vsys 1.

Select **Monitor** > **Logs** > **Traffic** to view traffic logs. Each entry has the option to display columns with information necessary to configure QoS in a virtual system environment:

- virtual system
- egress interface
- ingress interface
- source zone
- destination zone

To display a column if it is not displayed by default:

- Click any column header to add a column to the log:



- Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface, as well as source and destination zones, in the **Source** and **Destination** sections:



For example, for web-browsing traffic from VSYS 1, the ingress interface is ethernet 1/2, the egress interface is ethernet 1/1, the source zone is trust and the destination zone is untrust.

STEP 4 | Create a QoS Profile.

© 2018 Palo Alto Networks, Inc.

You can edit any existing QoS Profile, including the default, by clicking the profile name.

1. Select **Network** > **Network Profiles** > **QoS Profile** and click **Add** to open the QoS Profile dialog.
2. Enter a descriptive **Profile Name**.
3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS profile.
4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS profile.

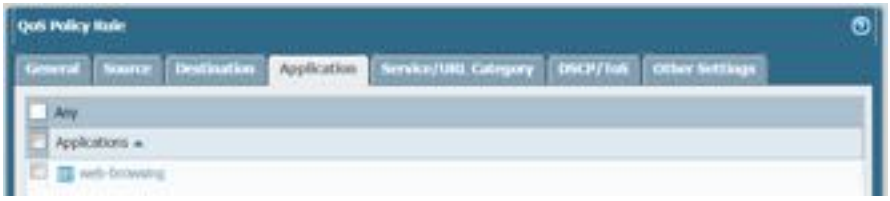
   Any traffic that exceeds the QoS profile's egress guaranteed limit is best effort but is not guaranteed.
5. In the Classes section of the **QoS Profile**, specify how to treat up to eight individual QoS classes:
   1. Click **Add** to add a class to the QoS Profile.
   2. Select the **Priority** for the class.
   3. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.
   4. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.
6. Click **OK** to save the QoS profile.

STEP 5 | Create a QoS policy.

In an environment with multiple virtual systems, traffic spans more than one virtual system. Because of this, when you are enabling QoS for a virtual system, you must define traffic to receive QoS treatment based on source and destination zones. This ensures that the traffic is prioritized and shaped only for that virtual system (and not for other virtual systems through which the traffic might flow).

1. Select **Policies** > **QoS** and **Add** a QoS Policy Rule.
2. Select **General** and give the QoS Policy Rule a descriptive **Name**.
3. Specify the traffic to which the QoS policy rule will apply. Use the **Source**, **Destination**, **Application**, and **Service/URL Category** tabs to define matching parameters for identifying traffic.

   For example, select **Application** and **Add** web-browsing to apply the QoS policy rule to that application:



4. Select **Source** and **Add** the source zone of vsys 1 web-browsing traffic.

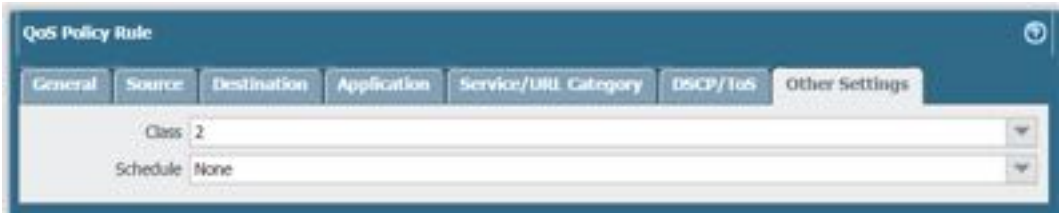


5. Select **Destination** and **Add** the destination zone of vsys 1 web-browsing traffic.



© 2018 Palo Alto Networks, Inc.

6. Select **Other Settings** and select a **QoS Class** to assign to the QoS policy rule. For example, assign Class 2 to web-browsing traffic on vsys 1:



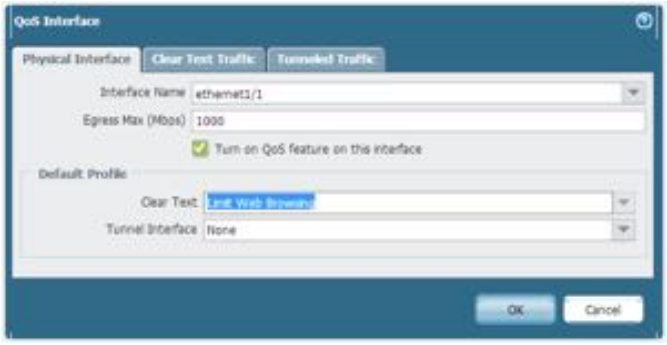
7. Click **OK** to save the QoS policy rule.

STEP 6 | Enable the QoS Profile on a physical interface.

It is a best practice to always define the Egress Max value for a QoS interface.

1. Select **Network** > **QoS** and click **Add** to open the QoS Interface dialog.
2. Enable QoS on the physical interface:
   1. On the **Physical Interface** tab, select the **Interface Name** of the interface to apply the QoS Profile to.

      In this example, ethernet 1/1 is the egress interface for web-browsing traffic on vsys 1 (see Step 2).



   2. Select **Turn on QoS feature on this interface**.
3. On the **Physical Interface** tab, select the default QoS profile to apply to all **Clear Text** traffic.

   (Optional) Use the **Tunnel Interface** field to apply a default QoS profile to all tunneled traffic.
4. (Optional) On the **Clear Text Traffic** tab, configure additional QoS settings for clear text traffic:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
   - Click **Add** to apply a QoS Profile to selected clear text traffic, further selecting the traffic for QoS treatment according to source interface and source subnet (creating a QoS node).
5. (Optional) On the **Tunneled Traffic** tab, configure additional QoS settings for tunnel interfaces:
   - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
   - Click **Add** to associate a selected tunnel interface with a QoS Profile.
6. Click **OK** to save changes.
7. **Commit** the changes.

STEP 7 | Verify QoS configuration.

© 2018 Palo Alto Networks, Inc.

- Select **Network** > **QoS** to view the QoS Policies page. The **QoS Policies** page verifies that QoS is enabled and includes a **Statistics** link. Click the Statistics link to view QoS bandwidth, active sessions of a selected QoS node or class, and active applications for the selected QoS node or class.
- In a multi-vsys environment, sessions cannot span multiple systems. Multiple sessions are created for one traffic flow if the traffic passes through more than one virtual system. To browse sessions running on the firewall and view applied QoS Rules and QoS Classes, select **Monitor** > **Session Browser**.

© 2018 Palo Alto Networks, Inc.

Enforce QoS Based on DSCP Classification

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 to mark a session with a DSCP value as session traffic exits the firewall. This enables all inbound and outbound traffic for a session can receive continuous QoS treatment as it flows through your 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 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 best effort service provides (though packets with an EF codepoint will continue to take precedence over those with an AF codepoint).
- **Class Selector (CS)**: 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**: Create a custom codepoint to match to traffic by entering a **Codepoint Name** and **Binary Value**.

For example, select the **Assured Forwarding (AF)** to ensure traffic marked with an AF codepoint value has higher priority for reliable delivery over applications marked to receive lower priority.Use the following steps 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.

STEP 1 | Perform the preliminary steps to Configure QoS.

STEP 2 | Define the traffic to receive QoS treatment based on DSCP value.

1. Select **Policies** > **QoS** and **Add** or modify an existing QoS rule and populate required fields.
2. Select **DSCP/ToS** and select **Codepoints**.
3. **Add** DSCP/ToS codepoints for which you want to enforce QoS.
4. Select the **Type** of DSCP/ToS marking for the QoS rule to match to traffic:



It is a best practice to use a single DSCP type to manage and prioritize your network traffic.

© 2018 Palo Alto Networks, Inc.

5. Match the QoS policy to traffic on a more granular scale by specifying the **Codepoint** value. For example, with Assured Forwarding (AF) selected as the **Type** of DSCP value for the policy to match, further specify an AF **Codepoint** value such as AF11.

   When Expedited Forwarding (EF) is selected as the Type of DSCP marking, a granular Codepoint value cannot be specified. The QoS policy rule matches to traffic marked with any EF codepoint value.
6. Select **Other Settings** and assign a **QoS Class** to traffic matched to the QoS rule. In this example, assign Class 1 to sessions where a DSCP marking of AF11 is detected for the first packet in the session.
7. Click **OK** to save the QoS rule.

STEP 3 | Define the QoS priority for traffic to receive when it is matched to a QoS rule based the DSCP marking detected at the beginning of a session.

1. Select **Network** > **Network Profiles** > **QoS Profile** and **Add** or modify an existing QoS profile. For details on profile options to set priority and bandwidth for traffic, see QoS Concepts and Configure QoS.
2. **Add** or modify a profile class. For example, because Step 2 showed steps to classify AF11 traffic as Class 1 traffic, you could add or modify a **class1** entry.
3. Select a **Priority** for the class of traffic, such as **high**.
4. Click **OK** to save the QoS Profile.

STEP 4 | Enable QoS on an interface.

Select **Network** > **QoS** and **Add** or modify an existing interface and **Turn on QoS feature on this interface**.

In this example, traffic with an AF11 DSCP marking is matched to the QoS rule and assigned Class 1. The QoS profile enabled on the interface enforces high priority treatment for Class 1 traffic as it egresses the firewall (the session *outbound* traffic).

STEP 5 | Enable DSCP Marking.

Mark return traffic with a DSCP value, enabling the inbound flow for a session to be marked with the same DSCP value detected for the outbound flow.

1. Select **Policies** > **Security** and **Add** or modify a security policy.
2. Select **Actions** and in the **QoS Marking** drop-down, choose **Follow-Client-to-Server-Flow**.
3. Click **OK** to save your changes.

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.

STEP 6 | Commit the configuration.

Commit your changes.

© 2018 Palo Alto Networks, Inc.

QoS Use Cases

The following use cases demonstrate how to use QoS in common scenarios:

- Use Case: QoS for a Single User on page 721
- Use Case: QoS for Voice and Video Applications on page 723

Use Case: QoS for a Single User

A CEO finds that during periods of high network usage, she is unable to access enterprise applications to respond effectively to critical business communications. The IT admin wants to ensure that all traffic to and from the CEO receives preferential treatment over other employee traffic so that she is guaranteed not only access to, but high performance of, critical network resources.

STEP 1 | The admin creates the QoS profile ***CEO_traffic*** to define how traffic originating from the CEO will be treated and shaped as it flows out of the company network:



The admin assigns a guaranteed bandwidth (**Egress Guaranteed**) of 50 Mbps to ensure that the CEO will have that amount that bandwidth guaranteed to her at all times (more than she would need to use), regardless of network congestion.

The admin continues by designating Class 1 traffic as high priority and sets the profile's maximum bandwidth usage (**Egress Max**) to 1000 Mbps, the same maximum bandwidth for the interface that the admin will enable QoS on. The admin is choosing to not restrict the CEO's bandwidth usage in any way.

It is a best practice to populate the Egress Max field for a QoS profile, even if the max bandwidth of the profile matches the max bandwidth of the interface. The QoS profile's max bandwidth should never exceed the max bandwidth of the interface you are planning to enable QoS on.

STEP 2 | The admin creates a QoS policy to identify the CEO's traffic (**Policies** > **QoS**) and assigns it the class that he defined in the QoS profile (see prior step). Because User-ID is configured, the admin uses the **Source** tab in the QoS policy to singularly identify the CEO's traffic by her company network username. (If User-ID is not configured, the administrator could **Add** the CEO's IP address under **Source Address**. See User-ID.):



© 2018 Palo Alto Networks, Inc.

The admin associates the CEO's traffic with Class 1 (**Other Settings** tab) and then continues to populate the remaining required policy fields; the admin gives the policy a descriptive **Name** (**General** tab) and selects **Any** for the **Source Zone** (**Source** tab) and **Destination Zone** (**Destination** tab):

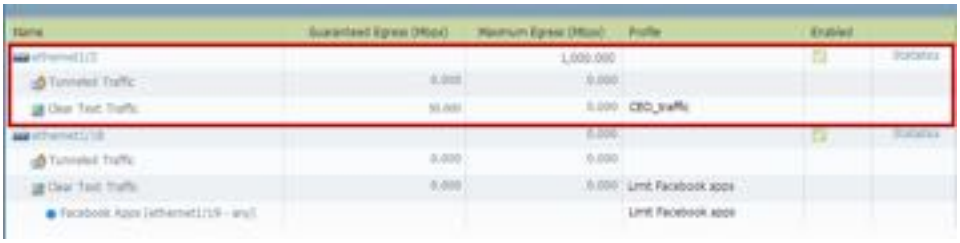


STEP 3 | Now that Class 1 is associated with the CEO's traffic, the admin enables QoS by checking **Turn on QoS feature on interface** and selecting the traffic flow's egress interface. The egress interface for the CEO's traffic flow is the external-facing interface, in this case, ethernet 1/2:



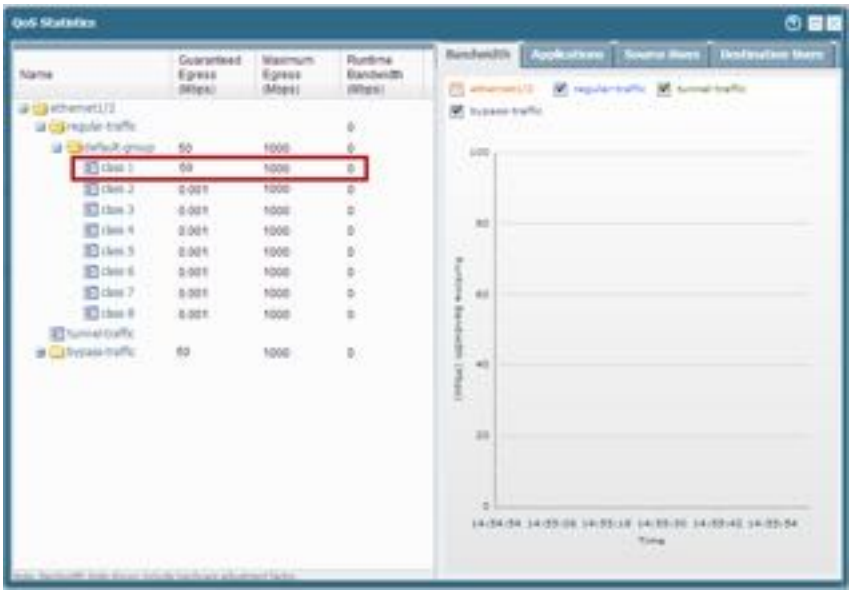
Because the admin wants to ensure that all traffic originating from the CEO is guaranteed by the QoS profile and associated QoS policy he created, he selects the *CEO_traffic* to apply to **Clear Text** traffic flowing from ethernet 1/2.

STEP 4 | After committing the QoS configuration, the admin navigates to the **Network** > **QoS** page to confirm that the QoS profile CEO_traffic is enabled on the external-facing interface, ethernet 1/2:



STEP 5 | He clicks **Statistics** to view how traffic originating with the CEO (Class 1) is being shaped as it flows from ethernet 1/2:

© 2018 Palo Alto Networks, Inc.



This case demonstrates how to apply QoS to traffic originating from a single source user. However, if you also wanted to guarantee or shape traffic to a destination user, you could configure a similar QoS setup. Instead of, or in addition to this work flow, create a QoS policy that specifies the user's IP address as the Destination Address on the Policies > QoS page (instead of specifying the user's source information) and then enable QoS on the network's internal-facing interface on the Network > QoS page (instead of the external-facing interface).

Use Case: QoS for Voice and Video Applications

Voice and video traffic is particularly sensitive to 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.

In this example, employees at a company branch office are experiencing difficulties and unreliability in using video conferencing and Voice over IP (VoIP) technologies to conduct business communications with other branch offices, with partners, and with customers. An IT admin intends to implement QoS in order to address these issues and ensure effective and reliable business communication for the branch employees. Because the admin wants to guarantee QoS to both incoming and outgoing network traffic, he will enable QoS on both the firewall's internal- and external-facing interfaces.

STEP 1 | The admin creates a QoS profile, defining Class 2 so that Class 2 traffic receives real-time priority and on an interface with a maximum bandwidth of 1000 Mbps, is guaranteed a bandwidth of 250 Mbps at all times, including peak periods of network usage.

Real-time priority is typically recommended for applications affected by latency, and is particularly useful in guaranteeing performance and quality of voice and video applications.

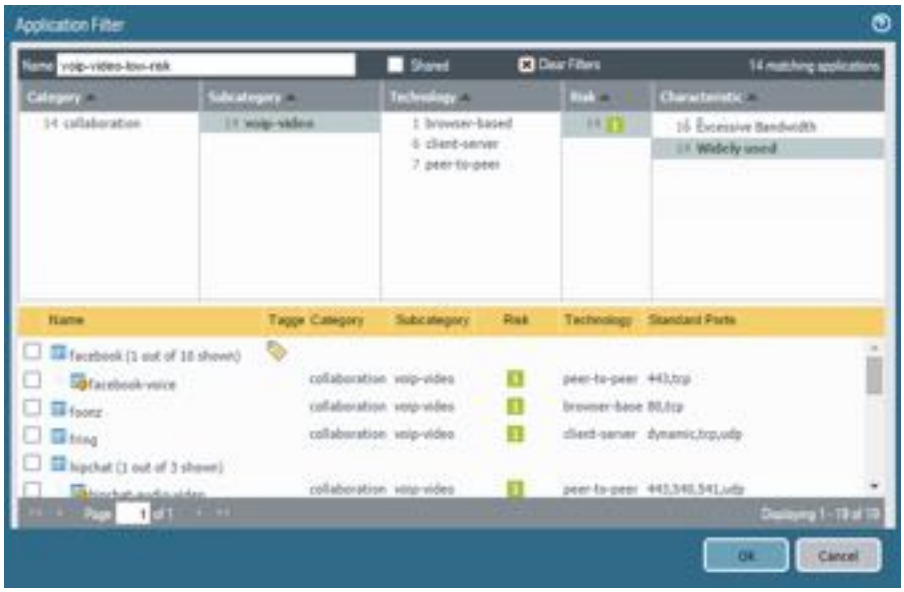
On the firewall web interface, the admin selects **Network** > **Network Profiles** > **Qos Profile** page, clicks **Add**, enters the **Profile Name** ensure voip-video traffic and defines Class 2 traffic.

© 2018 Palo Alto Networks, Inc.

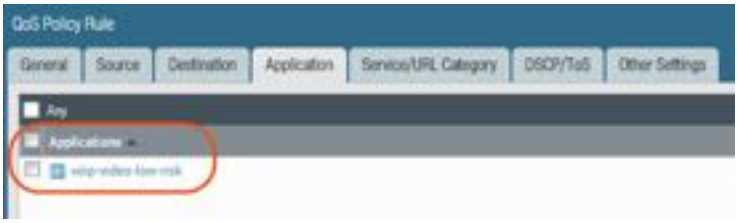


STEP 2 | The admin creates a QoS policy to identify voice and video traffic. Because the company does not have one standard voice and video application, the admin wants to ensure QoS is applied to a few applications that are widely and regularly used by employees to communicate with other offices, with partners, and with customers. On the **Policies** > **QoS** > **QoS Policy Rule** > **Applications** tab, the admin clicks **Add** and opens the **Application Filter** window. The admin continues by selecting criteria to filter the applications he wants to apply QoS to, choosing the Subcategory voip-video, and narrowing that down by specifying only voip-video applications that are both low-risk and widely-used.

The application filter is a dynamic tool that, when used to filter applications in the QoS policy, allows QoS to be applied to all applications that meet the criteria of voip-video, low risk, and widely used at any given time.



The admin names the **Application Filter** voip-video-low-risk and includes it in the QoS policy:



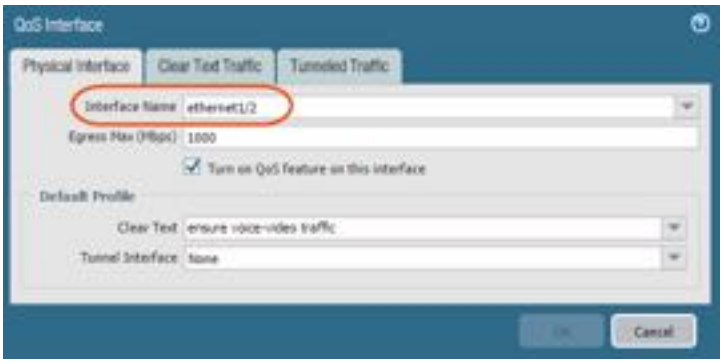
© 2018 Palo Alto Networks, Inc.

The admin names the QoS policy Voice-Video and selects Other Settings to assign all traffic matched to the policy Class 2. He is going to use the Voice-Video QoS policy for both incoming and outgoing QoS traffic, so he sets **Source** and **Destination** information to **Any**:



STEP 3 | Because the admin wants to ensure QoS for both incoming and outgoing voice and video communications, he enables QoS on the network's external-facing interface (to apply QoS to outgoing communications) and to the internal-facing interface (to apply QoS to incoming communications).

The admin begins by enabling the QoS profile he created, ensure voice-video traffic (Class 2 in this profile is associated with policy, Voice-Video) on the external-facing interface, in this case, ethernet 1/2.

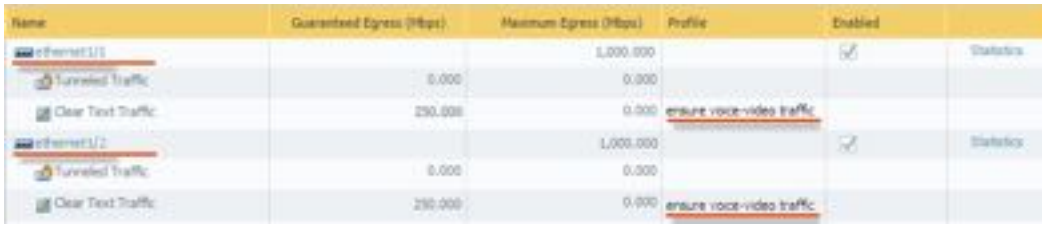


He then enables the same QoS profile ensure voip-video traffic on a second interface, the internal-facing interface (in this case, ethernet 1/1).



STEP 4 | The admin selects **Network** > **QoS** to confirm that QoS is enabled for both incoming and outgoing voice and video traffic:

© 2018 Palo Alto Networks, Inc.



The admin has successfully enabled QoS on both the network's internal- and external-facing interfaces. Real-time priority is now ensured for voice and video application traffic as it flows both into and out of the network, ensuring that these communications, which are particularly sensitive to latency and jitter, can be used reliably and effectively to perform both internal and external business communications.

VPNs

Virtual private networks (VPNs) create tunnels that allow users/systems to connect securely over a public network, as if they were connecting over a local area network (LAN). To set up a VPN tunnel, you need a pair of devices that can authenticate each other and encrypt the flow of information between them. The devices can be a pair of Palo Alto Networks firewalls, or a Palo Alto Networks firewall along with a VPN-capable device from another vendor.

> VPN Deployments on page 729
> Site-to-Site VPN Overview on page 730
> Site-to-Site VPN Concepts on page 731
> Set Up Site-to-Site VPN on page 738
> Site-to-Site VPN Quick Configs on page 752

© 2018 Palo Alto Networks, Inc.

VPN Deployments

The Palo Alto Networks firewall supports the following VPN deployments:

- **Site-to-Site VPN**— A simple VPN that connects a central site and a remote site, or a hub and spoke VPN that connects a central site with multiple remote sites. The firewall uses the IP Security (IPSec) set of protocols to set up a secure tunnel for the traffic between the two sites. See Site-to-Site VPN Overview.
- **Remote User-to-Site VPN**—A solution that uses the GlobalProtect agent to allow a remote user to establish a secure connection through the firewall. This solution uses SSL and IPSec to establish a secure connection between the user and the site. Refer to the GlobalProtect Administrator's Guide.
- **Large Scale VPN**— The Palo Alto Networks GlobalProtect Large Scale VPN (LSVPN) provides a simplified mechanism to roll out a scalable hub and spoke VPN with up to 1,024 satellite offices. The solution requires Palo Alto Networks firewalls to be deployed at the hub and at every spoke. It uses certificates for device authentication, SSL for securing communication between all components, and IPSec to secure data. See Large Scale VPN (LSVPN).

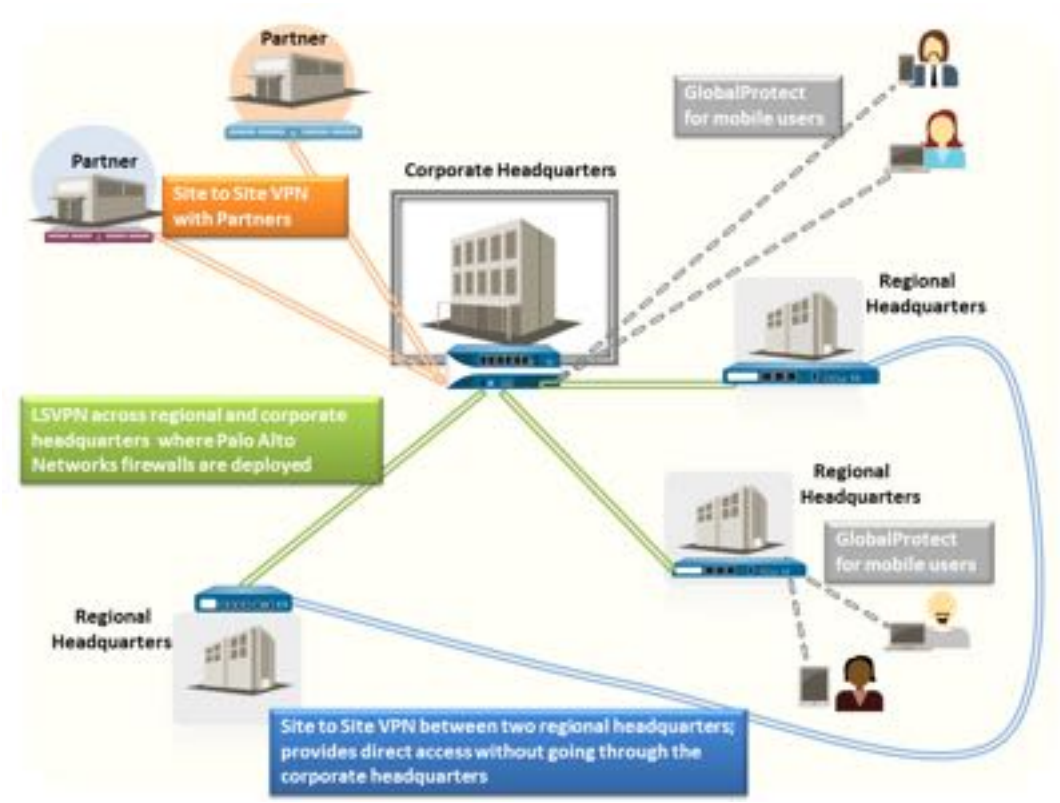


Figure 9: VPN Deployments

© 2018 Palo Alto Networks, Inc.

Site-to-Site VPN Overview

A VPN connection that allows you to connect two Local Area Networks (LANs) is called a site-to-site VPN. You can configure route-based VPNs to connect Palo Alto Networks firewalls located at two sites or to connect a Palo Alto Networks firewall with a third-party security device at another location. The firewall can also interoperate with third-party policy-based VPN devices; the Palo Alto Networks firewall supports route-based VPN.

The Palo Alto Networks firewall sets up a route-based VPN, where the firewall makes a routing decision based on the destination IP address. If traffic is routed to a specific destination through a VPN tunnel, then it is handled as VPN traffic.

The IP Security (IPSec) set of protocols is used to set up a secure tunnel for the VPN traffic, and the information in the TCP/IP packet is secured (and encrypted if the tunnel type is ESP). The IP packet (header and payload) is embedded in another IP payload, and a new header is applied and then sent through the IPSec tunnel. The source IP address in the new header is that of the local VPN peer and the destination IP address is that of the VPN peer on the far end of the tunnel. When the packet reaches the remote VPN peer (the firewall at the far end of the tunnel), the outer header is removed and the original packet is sent to its destination.

In order to set up the VPN tunnel, first the peers need to be authenticated. After successful authentication, the peers negotiate the encryption mechanism and algorithms to secure the communication. The Internet Key Exchange (IKE) process is used to authenticate the VPN peers, and IPSec Security Associations (SAs) are defined at each end of the tunnel to secure the VPN communication. IKE uses digital certificates or preshared keys, and the Diffie Hellman keys to set up the SAs for the IPSec tunnel.The SAs specify all of the parameters that are required for secure transmission— including the security parameter index (SPI), security protocol, cryptographic keys, and the destination IP address—encryption, data authentication, data integrity, and endpoint authentication.

The following figure shows a VPN tunnel between two sites. When a client that is secured by VPN Peer A needs content from a server located at the other site, VPN Peer A initiates a connection request to VPN Peer B. If the security policy permits the connection, VPN Peer A uses the IKE Crypto profile parameters (IKE phase 1) to establish a secure connection and authenticate VPN Peer B. Then, VPN Peer A establishes the VPN tunnel using the IPSec Crypto profile, which defines the IKE phase 2 parameters to allow the secure transfer of data between the two sites.



Figure 10: Site-to-Site VPN

© 2018 Palo Alto Networks, Inc.

Site-to-Site VPN Concepts

A VPN connection provides secure access to information between two or more sites. In order to provide secure access to resources and reliable connectivity, a VPN connection needs the following components:

- IKE Gateway on page 731
- Tunnel Interface on page 731
- Tunnel Monitoring on page 732
- Internet Key Exchange (IKE) for VPN on page 732
- IKEv2 on page 734

IKE Gateway

The Palo Alto Networks firewalls or a firewall and another security device that initiate and terminate VPN connections across the two networks are called the IKE Gateways. To set up the VPN tunnel and send traffic between the IKE Gateways, each peer must have an IP address—static or dynamic—or FQDN. The VPN peers use preshared keys or certificates to mutually authenticate each other.

The peers must also negotiate the mode—main or aggressive—for setting up the VPN tunnel and the SA lifetime in IKE Phase 1. Main mode protects the identity of the peers and is more secure because more packets are exchanged when setting up the tunnel. Main mode is the recommended mode for IKE negotiation if both peers support it. Aggressive mode uses fewer packets to set up the VPN tunnel and is hence faster but a less secure option for setting up the VPN tunnel.

See Set Up an IKE Gateway on page 738 for configuration details.

Tunnel Interface

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, for example the untrust zone. While the tunnel interface can be in the same security zone as the physical interface, for added security and better visibility, you can create a separate zone for the tunnel interface. If you create a separate zone for the tunnel interface, say a 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 in order 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 towards the IPSec VPN tunnel capacity of the firewall, and the tunnel capacity varies by the firewall model.

© 2018 Palo Alto Networks, Inc.

See Set Up an IPSec Tunnel on page 745 for configuration details.

Tunnel Monitoring

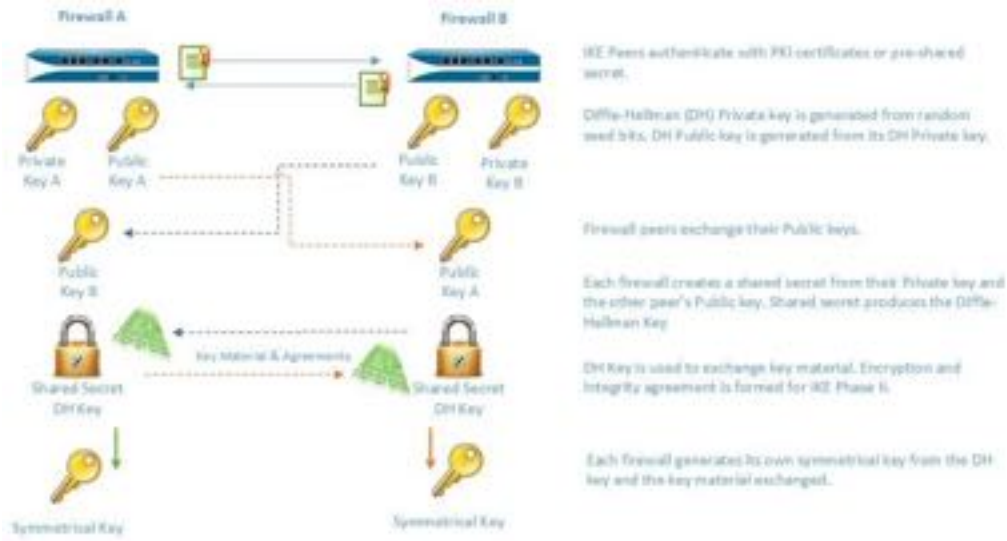
For a VPN tunnel, you can check connectivity to a destination IP address across the tunnel. The network monitoring profile on the firewall allows you to verify connectivity (using ICMP) to a destination IP address or a next hop at a specified polling interval, and to specify an action on failure to access the monitored IP address.

If the destination IP is unreachable, you 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.

See Set Up Tunnel Monitoring on page 747 for configuration details.

Internet Key Exchange (IKE) for VPN

The IKE process allows the VPN peers at both ends of the tunnel to encrypt and decrypt packets using mutually agreed-upon keys or certificate and method of encryption. The IKE process occurs in two phases: IKE Phase 1 and IKE Phase 2. Each of these phases use keys and encryption algorithms that are defined using cryptographic profiles— IKE crypto profile and IPSec crypto profile—and the result of the IKE negotiation is a Security Association (SA). An SA is a set of mutually agreed-upon keys and algorithms that are used by both VPN peers to allow the flow of data across the VPN tunnel. The following illustration depicts the key exchange process for setting up the VPN tunnel:



IKE Phase 1

In this phase, the firewalls use the parameters defined in the IKE Gateway configuration and the IKE Crypto profile to authenticate each other and set up a secure control channel. IKE Phase supports the use of preshared keys or digital certificates (which use public key infrastructure, PKI) for mutual authentication of the VPN peers. Preshared keys are a simple solution for securing smaller networks because they do not require the support of a PKI infrastructure. Digital certificates can be more convenient for larger networks or implementations that require stronger authentication security.

When using certificates, make sure that the CA issuing the certificate is trusted by both gateway peers and that the maximum length of certificates in the certificate chain is 5 or less. With IKE fragmentation enabled, the firewall can reassemble IKE messages with up to 5 certificates in the certificate chain and successfully establish a VPN tunnel.

© 2018 Palo Alto Networks, Inc.

The IKE Crypto profile defines the following options that are used in the IKE SA negotiation:

- Diffie-Hellman (DH) group for generating symmetrical keys for IKE.

  The Diffie-Hellman algorithm uses the private key of one party and the public key of the other to create a shared secret, which is an encrypted key that both VPN tunnel peers share. The DH groups supported on the firewall are: Group 1—768 bits, Group 2—1024 bits (default), Group 5—1536 bits, Group 14—2048 bits, Group 19—256-bit elliptic curve group, and Group 20—384-bit elliptic curve group.
- Authentication algorithms—sha1, sha 256, sha 384, sha 512, or md5
- Encryption algorithms—3des, aes-128-cbc, aes-192-cbc, aes-256-cbc, or des

IKE Phase 2

After the tunnel is secured and authenticated, in Phase 2 the channel is further secured for the transfer of data between the networks. IKE Phase 2 uses the keys that were established in Phase 1 of the process and the IPSec Crypto profile, which defines the IPSec protocols and keys used for the SA in IKE Phase 2.

The IPSEC uses the following protocols to enable secure communication:

- Encapsulating Security Payload (ESP)—Allows you to encrypt the entire IP packet, and authenticate the source and verify integrity of the data. While ESP requires that you encrypt and authenticate the packet, you can choose to only encrypt or only authenticate by setting the encryption option to Null; using encryption without authentication is discouraged.
- Authentication Header (AH)—Authenticates the source of the packet and verifies data integrity. AH does not encrypt the data payload and is unsuited for deployments where data privacy is important. AH is commonly used when the main concern is to verify the legitimacy of the peer, and data privacy is not required.

Table 6: Algorithms Supported for IPSEC Authentication and Encryption

ESP	AH
Diffie Hellman (DH) exchange options supported	
• Group 1—768 bits • Group 2—1024 bits (the default) • Group 5—1536 bits • Group 14—2048 bits. • Group 19— 256-bit elliptic curve group • Group 20—384-bit elliptic curve group • no-pfs—By default, perfect forward secrecy (PFS) is enabled, which means a new DH key is generated in IKE phase 2 using one of the groups listed above. This key is independent of the keys exchanged in IKE phase1 and provides better data transfer security. If you select no-pfs, the DH key created at phase 1 is not renewed and a single key is used for the IPSec SA negotiations. Both VPN peers must be enabled or disabled for PFS.	
Encryption algorithms supported	
• 3des	Triple Data Encryption Standard (3DES) with a security strength of 112 bits
• aes-128-cbc	Advanced Encryption Standard (AES) using cipher block chaining (CBC) with a security strength of 128 bits
• aes-192-cbc	AES using CBC with a security strength of 192 bits

© 2018 Palo Alto Networks, Inc.

ESP	AH
• aes-256-cbc	AES using CBC with a security strength of 256 bits
• aes-128-ccm	AES using Counter with CBC-MAC (CCM) with a security strength of 128 bits
• aes-128-gcm	AES using Galois/Counter Mode (GCM) with a security strength of 128 bits
• aes-256-gcm	AES using GCM with a security strength of 256 bits
• des	Data Encryption Standard (DES) with a security strength of 56 bits
Authentication algorithms supported	
• md5	• md5
• sha 1	• sha 1
• sha 256	• sha 256
• sha 384	• sha 384
• sha512	• sha 512

Methods of Securing IPSec VPN Tunnels (IKE Phase 2)

IPSec VPN tunnels can be secured using manual keys or auto keys. In addition, IPSec configuration options include Diffie-Hellman Group for key agreement, and/or an encryption algorithm and a hash for message authentication.

- **Manual Key**—Manual key is typically used if the Palo Alto Networks firewall is establishing a VPN tunnel with a legacy device, or if you want to reduce the overhead of generating session keys. If using manual keys, the same key must be configured on both peers.

  Manual keys are not recommended for establishing a VPN tunnel because the session keys can be compromised when relaying the key information between the peers; if the keys are compromised, the data transfer is no longer secure.
- **Auto Key**— Auto Key allows you to automatically generate keys for setting up and maintaining the IPSec tunnel based on the algorithms defined in the IPSec Crypto profile.

IKEv2

An IPSec VPN gateway uses IKEv1 or IKEv2 to negotiate the IKE security association (SA) and IPSec tunnel. IKEv2 is defined in RFC 5996.

Unlike IKEv1, which uses Phase 1 SA and Phase 2 SA, IKEv2 uses a child SA for Encapsulating Security Payload (ESP) or Authentication Header (AH), which is set up with an IKE SA.

NAT traversal (NAT-T) must be enabled on both gateways if you have NAT occurring on a device that sits between the two gateways. A gateway can see only the public (globally routable) IP address of the NAT device.

IKEv2 provides the following benefits over IKEv1:

© 2018 Palo Alto Networks, Inc.

- Tunnel endpoints exchange fewer messages to establish a tunnel. IKEv2 uses four messages; IKEv1 uses either nine messages (in main mode) or six messages (in aggressive mode).
- Built-in NAT-T functionality improves compatibility between vendors.
- Built-in health check automatically re-establishes a tunnel if it goes down. The liveness check replaces the Dead Peer Detection used in IKEv1.
- Supports traffic selectors (one per exchange). The traffic selectors are used in IKE negotiations to control what traffic can access the tunnel.
- Supports Hash and URL certificate exchange to reduce fragmentation.
- Resiliency against DoS attacks with improved peer validation. An excessive number of half-open SAs can trigger cookie validation.

Before configuring IKEv2, you should be familiar with the following concepts:

- Liveness Check on page 735
- Cookie Activation Threshold and Strict Cookie Validation on page 735
- Traffic Selectors on page 736
- Hash and URL Certificate Exchange on page 736
- SA Key Lifetime and Re-Authentication Interval on page 737

After you Set Up an IKE Gateway on page 738, if you chose IKEv2, perform the following optional tasks related to IKEv2 as required by your environment:

- Export a Certificate for a Peer to Access Using Hash and URL on page 740
- Import a Certificate for IKEv2 Gateway Authentication on page 741
- Change the Key Lifetime or Authentication Interval for IKEv2 on page 741
- Change the Cookie Activation Threshold for IKEv2 on page 742
- Configure IKEv2 Traffic Selectors on page 742

Liveness Check

The liveness check for IKEv2 is similar to Dead Peer Detection (DPD), which IKEv1 uses as the way to determine whether a peer is still available.

In IKEv2, the liveness check is achieved by any IKEv2 packet transmission or an empty informational message that the gateway sends to the peer at a configurable interval, five seconds by default. If necessary, the sender attempts the retransmission up to ten times. If it doesn't get a response, the sender closes and deletes the IKE_SA and corresponding CHILD_SAs. The sender will start over by sending out another IKE_SA_INIT message.

Cookie Activation Threshold and Strict Cookie Validation

Cookie validation is always enabled for IKEv2; it helps protect against half-SA DoS attacks. You can configure the global threshold number of half-open SAs that will trigger cookie validation. You can also configure individual IKE gateways to enforce cookie validation for every new IKEv2 SA.

- The **Cookie Activation Threshold** is a global VPN session setting that limits the number of simultaneous half-opened IKE SAs (default is 500). When the number of half-opened IKE SAs exceeds the **Cookie Activation Threshold**, the Responder will request a cookie, and the Initiator must respond with an IKE_SA_INIT containing a cookie to validate the connection. If the cookie validation is successful, another SA can be initiated. A value of 0 means that cookie validation is always on.

  The Responder does not maintain a state of the Initiator, nor does it perform a Diffie-Hellman key exchange, until the Initiator returns the cookie. IKEv2 cookie validation mitigates a DoS attack that would try to leave numerous connections half open.

  The **Cookie Activation Threshold** must be lower than the **Maximum Half Opened SA** setting. If you Change the Cookie Activation Threshold for IKEv2 on page 742 to a very high number (for example,

© 2018 Palo Alto Networks, Inc.

65534) and the **Maximum Half Opened SA** setting remained at the default value of 65535, cookie validation is essentially disabled.
- You can enable **Strict Cookie Validation** if you want cookie validation performed for every new IKEv2 SA a gateway receives, regardless of the global threshold. **Strict Cookie Validation** affects only the IKE gateway being configured and is disabled by default. With **Strict Cookie Validation** disabled, the system uses the **Cookie Activation Threshold** to determine whether a cookie is needed or not.

Traffic Selectors

In IKEv1, a firewall that has a route-based VPN needs to use a local and remote Proxy ID in order to set up an IPSec tunnel. Each peer compares its Proxy IDs with what it received in the packet in order to successfully negotiate IKE Phase 2. IKE Phase 2 is about negotiating the SAs to set up an IPSec tunnel. (For more information on Proxy IDs, see Tunnel Interface on page 731.)

In IKEv2, you can Configure IKEv2 Traffic Selectors on page 742, which are components of network traffic that are used during IKE negotiation. Traffic selectors are used during the CHILD_SA (tunnel creation) Phase 2 to set up the tunnel and to determine what traffic is allowed through the tunnel. The two IKE gateway peers must negotiate and agree on their traffic selectors; otherwise, one side narrows its address range to reach agreement. One IKE connection can have multiple tunnels; for example, you can assign different tunnels to each department to isolate their traffic. Separation of traffic also allows features such as QoS to be implemented.

The IPv4 and IPv6 traffic selectors are:

- **Source IP address**—A network prefix, address range, specific host, or wildcard.
- **Destination IP address**—A network prefix, address range, specific host, or wildcard.
- **Protocol**—A transport protocol, such as TCP or UDP.
- **Source port**—The port where the packet originated.
- **Destination port**—The port the packet is destined for.

During IKE negotiation, there can be multiple traffic selectors for different networks and protocols. For example, the Initiator might indicate that it wants to send TCP packets from 172.168.0.0/16 through the tunnel to its peer, destined for 198.5.0.0/16. It also wants to send UDP packets from 172.17.0.0/16 through the same tunnel to the same gateway, destined for 0.0.0.0 (any network). The peer gateway must agree to these traffic selectors so that it knows what to expect.

It is possible that one gateway will start negotiation using a traffic selector that is a more specific IP address than the IP address of the other gateway.

- For example, gateway A offers a source IP address of 172.16.0.0/16 and a destination IP address of 192.16.0.0/16. But gateway B is configured with 0.0.0.0 (any source) as the source IP address and 0.0.0.0 (any destination) as the destination IP address. Therefore, gateway B narrows down its source IP address to 192.16.0.0/16 and its destination address to 172.16.0.0/16. Thus, the narrowing down accommodates the addresses of gateway A and the traffic selectors of the two gateways are in agreement.
- If gateway B (configured with source IP address 0.0.0.0) is the Initiator instead of the Responder, gateway A will respond with its more specific IP addresses, and gateway B will narrow down its addresses to reach agreement.

Hash and URL Certificate Exchange

IKEv2 supports Hash and URL Certificate Exchange, which is used during an IKEv2 negotiation of an SA. You store the certificate on an HTTP server, which is specified by a URL. The peer fetches the certificate from the server based on receiving the URL to the server. The hash is used to check whether the content of the certificate is valid or not. Thus, the two peers exchange certificates with the HTTP CA rather than with each other.

© 2018 Palo Alto Networks, Inc.

The hash part of Hash and URL reduces the message size and thus Hash and URL is a way to reduce the likelihood of packet fragmentation during IKE negotiation. The peer receives the certificate and hash that it expects, and thus IKE Phase 1 has validated the peer. Reducing fragmentation occurrences helps protect against DoS attacks.

You can enable the Hash and URL certificate exchange when configuring an IKE gateway by selecting **HTTP Certificate Exchange** and entering the **Certificate URL**. The peer must also use Hash and URL certificate exchange in order for the exchange to be successful. If the peer cannot use Hash and URL, X.509 certificates are exchanged similarly to how they are exchanged in IKEv1.

If you enable the Hash and URL certificate exchange, you must export your certificate to the certificate server if it is not already there. When you export the certificate, the file format should be **Binary Encoded Certificate (DER)**. See Export a Certificate for a Peer to Access Using Hash and URL on page 740.

SA Key Lifetime and Re-Authentication Interval

In IKEv2, two IKE crypto profile values, **Key Lifetime** and **IKEv2 Authentication Multiple**, control the establishment of IKEv2 IKE SAs. The key lifetime is the length of time that a negotiated IKE SA key is effective. Before the key lifetime expires, the SA must be re-keyed; otherwise, upon expiration, the SA must begin a new IKEv2 IKE SA re-key. The default value is 8 hours.

The re-authentication interval is derived by multiplying the **Key Lifetime** by the **IKEv2 Authentication Multiple**. The authentication multiple defaults to 0, which disables the re-authentication feature.

The range of the authentication multiple is 0-50. So, if you were to configure an authentication multiple of 20, for example, the system would perform re-authentication every 20 re-keys, which is every 160 hours. That means the gateway could perform Child SA creation for 160 hours before the gateway must re-authenticate with IKE to recreate the IKE SA from scratch.

In IKEv2, the Initiator and Responder gateways have their own key lifetime value, and the gateway with the shorter key lifetime is the one that will request that the SA be re-keyed.

© 2018 Palo Alto Networks, Inc.

Set Up Site-to-Site VPN

To set up site-to-site VPN:

- Make sure that your Ethernet interfaces, virtual routers, and zones are configured properly. For more information, see Configure Interfaces and Zones.
- Create your tunnel interfaces. Ideally, put the tunnel interfaces in a separate zone, so that tunneled traffic can use different policies.
- Set up static routes or assign routing protocols to redirect traffic to the VPN tunnels. To support dynamic routing (OSPF, BGP, RIP are supported), you must assign an IP address to the tunnel interface.
- Define IKE gateways for establishing communication between the peers across each end of the VPN tunnel; also define the cryptographic profile that specifies the protocols and algorithms for identification, authentication, and encryption to be used for setting up VPN tunnels in IKEv1 Phase 1. See Set Up an IKE Gateway and Define IKE Crypto Profiles.
- Configure the parameters that are needed to establish the IPSec connection for transfer of data across the VPN tunnel; See Set Up an IPSec Tunnel. For IKEv1 Phase-2, see Define IPSec Crypto Profiles.
- (Optional) Specify how the firewall will monitor the IPSec tunnels. See Set Up Tunnel Monitoring.
- Define security policies to filter and inspect the traffic.



If there is a deny rule at the end of the security rulebase, intra-zone traffic is blocked unless otherwise allowed. Rules to allow IKE and IPSec applications must be explicitly included above the deny rule.



If your VPN traffic is passing through (not originating or terminating on) a PA-7000 Series or PA-5200 Series firewall, configure bi-directional Security policy rules to allow the ESP or AH traffic in both directions.

When these tasks are complete, the tunnel is ready for use. Traffic destined for the zones/addresses defined in policy is automatically routed properly based on the destination route in the routing table, and handled as VPN traffic. For a few examples on site-to-site VPN, see Site-to-Site VPN Quick Configs.

For troubleshooting purposes, you can Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel.

Set Up an IKE Gateway

To set up a VPN tunnel, the VPN peers or gateways must authenticate each other using preshared keys or digital certificates and establish a secure channel in which to negotiate the IPSec security association (SA) that will be used to secure traffic between the hosts on each side.

STEP 1 | Define the IKE Gateway.

1. Select **Network** > **Network Profiles** > **IKE Gateways**, click **Add**, and on the **General** tab, enter the **Name** of the gateway.
2. For **Version**, select **IKEv1 only mode**, **IKEv2 only mode**, or **IKEv2 preferred mode**. The IKE gateway begins its negotiation with its peer in the mode specified here. If you select **IKEv2 preferred mode**, the two peers will use IKEv2 if the remote peer supports it; otherwise they will use IKEv1.

   The **Version** selection also determines which options are available on the **Advanced Options** tab.

STEP 2 | Establish the local endpoint of the tunnel (gateway).

1. For **Address Type**, click **IPv4** or **IPv6**.
2. Select the physical, outgoing **Interface** on the firewall where the local gateway resides.

© 2018 Palo Alto Networks, Inc.

3. From the **Local IP Address** drop-down, select the IP address that will be used as the endpoint for the VPN connection. This is the external-facing interface with a publicly routable IP address on the firewall.

STEP 3 | Establish the peer at the far end of the tunnel (gateway).

1. Select the **Peer IP Type** to be a **Static** or **Dynamic** address assignment.
2. If the **Peer IP Address** is static, enter the IP address of the peer.

STEP 4 | Specify how the peer is authenticated.

Select the **Authentication** method: **Pre-Shared Key** or **Certificate**. If you choose Pre-Shared Key, proceed to the next step. If you choose Certificate, skip to Configure certificate-based authentication.

STEP 5 | Configure a pre-shared key.

1. Enter a **Pre-shared Key**, which is the security key to use for authentication across the tunnel. Re-enter the value to **Confirm Pre-shared Key**. Use a maximum of 255 ASCII or non-ASCII characters.

   BEST PRACTICE: Generate a key that is difficult to crack with dictionary attacks; use a pre-shared key generator, if necessary.
2. For **Local Identification**, choose from the following types and enter a value that you determine: **FQDN (hostname)**, **IP address**, **KEYID (binary format ID string in HEX)**, **User FQDN (email address)**. Local identification defines the format and identification of the local gateway. If no value is specified, the local IP address will be used as the local identification value.
3. For **Peer Identification**, choose from the following types and enter the value: **FQDN (hostname)**, **IP address**, **KEYID (binary format ID string in HEX)**, **User FQDN (email address)**. Peer identification defines the format and identification of the peer gateway. If no value is specified, the peer IP address will be used as the peer identification value.
4. Proceed to Step 7 and continue from there.

STEP 6 | Configure certificate-based authentication.

Perform the remaining steps in this procedure if you selected **Certificate** as the method of authenticating the peer gateway at the opposite end of the tunnel.

1. Select a **Local Certificate** that is already on the firewall from the drop-down, or **Import** a certificate, or **Generate** to create a new certificate.
   - If you want to **Import** a certificate, Import a Certificate for IKEv2 Gateway Authentication and then return to this task.
   - If you want to **Generate** a new certificate, generate a certificate on the firewall and then return to this task.
2. Click the **HTTP Certificate Exchange** check box if you want to configure Hash and URL (IKEv2 only). For an HTTP certificate exchange, enter the **Certificate URL**. For more information, see Hash and URL Certificate Exchange.
3. Select the **Local Identification** type from the following: **Distinguished Name (Subject)**, **FQDN (hostname)**, **IP address**, **User FQDN (email address)**, and enter the value. Local identification defines the format and identification of the local gateway.
4. Select the **Peer Identification** type from the following: **Distinguished Name (Subject)**, **FQDN (hostname)**, **IP address**, **User FQDN (email address)**, and enter the value. Peer identification defines the format and identification of the peer gateway.
5. Select one type of **Peer ID Check**:
   - **Exact**—Check this to ensure that the local setting and peer IKE ID payload match exactly.
   - **Wildcard**—Check this to allow the peer identification to match as long as every character before the wildcard (*) matches. The characters after the wildcard need not match.

© 2018 Palo Alto Networks, Inc.

6. Click **Permit peer identification and certificate payload identification mismatch** if you want to allow a successful IKE SA even when the peer identification does not match the peer identification in the certificate.
7. Choose a **Certificate Profile** from the drop-down. A certificate profile contains information about how to authenticate the peer gateway.
8. Click **Enable strict validation of peer's extended key use** if you want to strictly control how the key can be used.

STEP 7 | Configure advanced options for the gateway.

1. Select the **Advanced Options** tab.
2. In the Common Options section, **Enable Passive Mode** if you want the firewall to only respond to IKE connection requests and never initiate them.
3. **Enable NAT Traversal** if you have a device performing NAT between the gateways, to have UDP encapsulation used on IKE and UDP protocols, enabling them to pass through intermediate NAT devices.
4. If you chose **IKEv1 only mode** earlier, on the IKEv1 tab:
   - Choose **auto**, **aggressive**, or **main** for the **Exchange Mode**. When a device is set to use **auto** exchange mode, it can accept both **main** mode and **aggressive** mode negotiation requests; however, whenever possible, it initiates negotiation and allows exchanges in **main** mode.

     If the exchange mode is not set to auto, you must configure both peers with the same exchange mode to allow each peer to accept negotiation requests.
   - Select an existing profile or keep the default profile from **IKE Crypto Profile** drop-down. For details on defining an IKE Crypto profile, see Define IKE Crypto Profiles.
   - (Only if using certificate-based authentication and the exchange mode is not set to aggressive mode) Click **Enable Fragmentation** to enable the firewall to operate with IKE Fragmentation.
   - Click **Dead Peer Detection** and enter an **Interval** (range is 2-100 seconds). For **Retry**, define the time to delay (range is 2-100 seconds) before attempting to re-check availability. Dead peer detection identifies inactive or unavailable IKE peers by sending an IKE phase 1 notification payload to the peer and waiting for an acknowledgment.
5. If you chose **IKEv2 only mode** or **IKEv2 preferred mode** in 1, on the IKEv2 tab:
   - Select an **IKE Crypto Profile** from the drop-down, which configures IKE Phase 1 options such as the DH group, hash algorithm, and ESP authentication. For information about IKE crypto profiles, see IKE Phase 1.
   - Enable **Strict Cookie Validation** if you want to always enforce cookie validation on IKEv2 SAs for this gateway. See Cookie Activation Threshold and Strict Cookie Validation.
   - **Enable Liveness Check** and enter an **Interval (sec)** (default is 5) if you want to have the gateway send a message request to its gateway peer, requesting a response. If necessary, the Initiator attempts the liveness check up to 10 times. If it doesn't get a response, the Initiator closes and deletes the IKE_SA and CHILD_SA. The Initiator will start over by sending out another IKE_SA_INIT.

STEP 8 | Commit your changes.

Click **OK** and **Commit**.

Export a Certificate for a Peer to Access Using Hash and URL

IKEv2 supports Hash and URL Certificate Exchange on page 736 as a method of having the peer at the remote end of the tunnel fetch the certificate from a server where you have exported the certificate. Perform this task to export your certificate to that server. You must have already created a certificate using **Device** > **Certificate Management**.

© 2018 Palo Alto Networks, Inc.

STEP 1 | Select **Device** > **Certificates**, and if your platform supports multiple virtual systems, for **Location**, select the appropriate virtual system.

STEP 2 | On the **Device Certificates** tab, select the certificate to **Export** to the server.



The status of the certificate should be valid, not expired. The firewall will not stop you from exporting an invalid certificate.

STEP 3 | For **File Format**, select **Binary Encoded Certificate (DER)**.

STEP 4 | Leave **Export private key** clear. Exporting the private key is unnecessary for Hash and URL.

STEP 5 | Click **OK**.

Import a Certificate for IKEv2 Gateway Authentication

Perform this task if you are authenticating a peer for an IKEv2 gateway and you did not use a local certificate already on the firewall; you want to import a certificate from elsewhere.

This task presumes that you selected **Network** > **IKE Gateways**, added a gateway, and for **Local Certificate**, you clicked **Import**.

STEP 1 | Import a certificate.

1. Select **Network** > **IKE Gateways**, **Add** a gateway, and on the **General** tab, for **Authentication**, select **Certificate**. For **Local Certificate**, click **Import**.
2. In the Import Certificate window, enter a **Certificate Name** for the certificate you are importing.
3. Select **Shared** if this certificate is to be shared among multiple virtual systems.
4. For **Certificate File**, **Browse** to the certificate file. Click on the file name and click **Open**, which populates the **Certificate File** field.
5. For **File Format**, select one of the following:
   - **Base64 Encoded Certificate (PEM)**—Contains the certificate, but not the key. It is cleartext.
   - **Encrypted Private Key and Certificate (PKCS12)**—Contains both the certificate and the key.
6. Select **Import private key** if the key is in a different file from the certificate file. The key is optional, with the following exception:
   - You must import a key if you set the **File Format** to **PEM**. Enter a **Key file** by clicking **Browse** and navigating to the key file to import.
   - Enter a **Passphrase** and **Confirm Passphrase**.
7. Click **OK**.

STEP 2 | Continue to the next task.

Step Configure certificate-based authentication.

Change the Key Lifetime or Authentication Interval for IKEv2

This task is optional; the default setting of the IKEv2 IKE SA re-key lifetime is 8 hours. The default setting of the IKEv2 Authentication Multiple is 0, meaning the re-authentication feature is disabled. For more information, see SA Key Lifetime and Re-Authentication Interval.

To change the default values, perform the following task. A prerequisite is that an IKE crypto profile already exists.

STEP 1 | Change the SA key lifetime or authentication interval for an IKE Crypto profile.

© 2018 Palo Alto Networks, Inc.

1. Select **Network** > **Network Profiles** > **IKE Crypto** and select the IKE Crypto profile that applies to the local gateway.
2. For the **Key Lifetime**, select a unit (**Seconds**, **Minutes**, **Hours**, or **Days**) and enter a value. The minimum is three minutes.
3. For **IKE Authentication Multiple**, enter a value, which is multiplied by the lifetime to determine the re-authentication interval.

STEP 2 | Commit your changes.

Click **OK** and **Commit**.

Change the Cookie Activation Threshold for IKEv2

Perform the following task if you want a firewall to have a threshold different from the default setting of 500 half-opened SA sessions before cookie validation is required. For more information about cookie validation, see Cookie Activation Threshold and Strict Cookie Validation.

STEP 1 | Change the Cookie Activation Threshold.

1. Select **Device** > **Setup** > **Session** and edit the VPN Session Settings. For **Cookie Activation Threshold**, enter the maximum number of half-opened SAs that are allowed before the responder requests a cookie from the initiator (range is 0-65,535; default is 500).
2. Click **OK**.

STEP 2 | Commit your changes.

Click **OK** and **Commit**.

Configure IKEv2 Traffic Selectors

In IKEv2, you can configure Traffic Selectors on page 736, which are components of network traffic that are used during IKE negotiation. Traffic selectors are used during the CHILD_SA (tunnel creation) Phase 2 to set up the tunnel and to determine what traffic is allowed through the tunnel. The two IKE gateway peers must negotiate and agree on their traffic selectors; otherwise, one side narrows its address range to reach agreement. One IKE connection can have multiple tunnels; for example, you can assign different tunnels to each department to isolate their traffic. Separation of traffic also allows features such as QoS to be implemented. Use the following workflow to configure traffic selectors.

STEP 1 | Select **Network** > **IPSec Tunnels** > **Proxy IDs**.

STEP 2 | Select the **IPv4** or **IPv6** tab.

STEP 3 | Click **Add** and enter the **Name** in the **Proxy ID** field.

STEP 4 | In the **Local** field, enter the **Source IP Address**.

STEP 5 | In the **Remote** field, enter the **Destination IP Address**.

STEP 6 | In the **Protocol** field, select the transport protocol (**TCP** or **UDP**) from the drop-down.

STEP 7 | Click **OK**.

© 2018 Palo Alto Networks, Inc.

Define Cryptographic Profiles

A cryptographic profile specifies the ciphers used for authentication and/or encryption between two IKE peers, and the lifetime of the key. The time period between each renegotiation is known as the lifetime; when the specified time expires, the firewall renegotiates a new set of keys.

For securing communication across the VPN tunnel, the firewall requires IKE and IPSec cryptographic profiles for completing IKE phase 1 and phase 2 negotiations, respectively. The firewall includes a default IKE crypto profile and a default IPSec crypto profile that are ready for use.

- Define IKE Crypto Profiles on page 743
- Define IPSec Crypto Profiles on page 744

Define IKE Crypto Profiles

The IKE crypto profile is used to set up the encryption and authentication algorithms used for the key exchange process in IKE Phase 1, and lifetime of the keys, which specifies how long the keys are valid. To invoke the profile, you must attach it to the IKE Gateway configuration.



All IKE gateways configured on the same interface or local IP address must use the same crypto profile.

STEP 1 | Create a new IKE profile.

1. Select **Network** > **Network Profiles** > **IKE Crypto** and select **Add**.
2. Enter a **Name** for the new profile.

STEP 2 | Specify the DH (Diffie–Hellman) Group for key exchange and the Authentication and Encryption algorithms.

Click **Add** in the corresponding sections (DH Group, Authentication, and Encryption) and select from the drop-downs.

If you are not certain of what the VPN peers support, add multiple groups or algorithms in the order of most-to-least secure as follows; the peers negotiate the strongest supported group or algorithm to establish the tunnel:

- DH Group—**group20**, **group19**, **group14**, **group5**, **group2**, and **group1**.
- Authentication—**sha512**, **sha384**, **sha256**, **sha1**, **md5**.
- Encryption—**aes-256-cbc**, **aes-192-cbc**, **aes-128-cbc**, **3des**, **des**.



As a best practice, choose the strongest authentication and encryption algorithms the peer can support. For the authentication algorithm, use SHA-256 or higher (SHA-384 or higher preferred for long-lived transactions). Do not use SHA-1 or MD5. For the encryption algorithm, use AES; DES and 3DES are weak and vulnerable.

STEP 3 | Specify the duration for which the key is valid and the re-authentication interval.

For details, see SA Key Lifetime and Re-Authentication Interval.

1. In the **Key Lifetime** fields, specify the period (in seconds, minutes, hours, or days) for which the key is valid (range is 3 minutes to 365 days; default is 8 hours). When the key expires, the firewall renegotiates a new key. A lifetime is the period between each renegotiation.
2. For the **IKEv2 Authentication Multiple**, specify a value (range is 0-50; default is 0) that is multiplied by the **Key Lifetime** to determine the authentication count. The default value of 0 disables the re-authentication feature.

STEP 4 | Commit your IKE Crypto profile.

© 2018 Palo Alto Networks, Inc.

Click **OK** and click **Commit**.

STEP 5 | Attach the IKE Crypto profile to the IKE Gateway configuration.

See Configure advanced options for the gateway.

Define IPSec Crypto Profiles

The IPSec crypto profile is invoked in IKE Phase 2. It specifies how the data is secured within the tunnel when Auto Key IKE is used to automatically generate keys for the IKE SAs.

STEP 1 | Create a new IPSec profile.

1. Select **Network** > **Network Profiles** > **IPSec Crypto** and select **Add**.
2. Enter a **Name** for the new profile.
3. Select the **IPSec Protocol**—ESP or AH—that you want to apply to secure the data as it traverses across the tunnel.

   As a best practice, select ESP (Encapsulating Security payload) over AH (Authentication Header) because ESP offers both confidentiality and authentication for the connection where as AH offers only authentication.
4. Click **Add** and select the **Authentication** and **Encryption** algorithms for ESP, and **Authentication** algorithms for AH, so that the IKE peers can negotiate the keys for the secure transfer of data across the tunnel.

   If you are not certain of what the IKE peers support, add multiple algorithms in the order of most-to-least secure as follows; the peers negotiate the strongest supported algorithm to establish the tunnel:

   - Encryption—**aes-256-gcm**, **aes-256-cbc**, **aes-192-cbc**, **aes-128-gcm**, **aes-128-ccm** (the VM-Series firewall doesn't support this option), **aes-128-cbc**, **3des**, **des**.
   - Authentication—**sha512**, **sha384**, **sha256**, **sha1**, **md5**.

   As a best practice, choose the strongest authentication and encryption algorithms the peer can support. For the authentication algorithm, use SHA-256 or higher (SHA-384 or higher preferred for long-lived transactions). Do not use SHA-512, SHA-1, or MD5. For the encryption algorithm, use AES; DES and 3DES are weak and vulnerable.

STEP 2 | Select the DH Group to use for the IPSec SA negotiations in IKE phase 2.

Select the key strength that you want to use from the **DH Group** drop-down.

If you are not certain of what the VPN peers support, add multiple groups in the order of most-to-least secure as follows; the peers negotiate the strongest supported group to establish the tunnel: **group20**, **group19**, **group14**, **group5**, **group2**, and **group1**.

Select **no-pfs** if you do not want to renew the key that was created at phase 1; the current key is reused for the IPSEC SA negotiations.

STEP 3 | Specify the duration of the key—time and volume of traffic.

Using a combination of time and traffic volume allows you to ensure safety of data.

Select the **Lifetime** or time period for which the key is valid in seconds, minutes, hours, or days (range is 3 minutes to 365 days). When the specified time expires, the firewall will renegotiate a new set of keys.

Select the **Lifesize** or volume of data after which the keys must be renegotiated.

STEP 4 | Commit your IPSec profile.

© 2018 Palo Alto Networks, Inc.

Click **OK** and click **Commit**.

STEP 5 | Attach the IPSec Profile to an IPSec tunnel configuration.

See Set up key exchange.

Set Up an IPSec Tunnel

The IPSec tunnel configuration allows you to authenticate and/or encrypt the data (IP packet) as it traverses across the tunnel.

If you are setting up the Palo Alto Networks 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 Palo Alto Networks firewall to work with a policy-based VPN peer, 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 Palo Alto Networks 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.

STEP 1 | Select **Network** > **IPSec Tunnels** and then **Add** a new tunnel configuration.

STEP 2 | On the General tab, enter a **Name** for the new tunnel.

STEP 3 | Select the **Tunnel interface** that will be used to set up the IPSec tunnel.

To create a new tunnel interface:

1. Select **Tunnel Interface** > **New Tunnel Interface**. (You can also select **Network** > **Interfaces** > **Tunnel** and click **Add**.)
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the Config tab, select the Security Zone drop-down to define the zone as follows:

Use your trust zone as the termination point for the tunnel—Select the zone from the drop-down. Associating the tunnel interface with the same zone (and virtual router) as the external-facing interface on which the packets enter the firewall mitigates the need to create inter-zone routing.

Or:

Create a separate zone for VPN tunnel termination (Recommended)—Select New Zone, define a Name for the new zone (for example vpn-corp), and click OK.

1. In the Virtual Router drop-down, select default.
2. (Optional) If you want to assign an IPv4 address to the tunnel interface, select the IPv4 tab, and Add the IP address and network mask, for example 10.31.32.1/32.
3. Click **OK**.

STEP 4 | (Optional) Enable IPv6 on the tunnel interface.

1. Select the IPv6 tab on **Network** > **Interfaces** > **Tunnel** > **IPv6**.
2. Select the check box to **Enable IPv6 on the interface**.

   This option allows you to route IPv6 traffic over an IPv4 IPSec tunnel and will provide confidentiality between IPv6 networks. The IPv6 traffic is encapsulated by IPv4 and then ESP. To route IPv6 traffic to the tunnel, you can use a static route to the tunnel, or use OSPFv3, or use a Policy-Based Forwarding (PBF) rule to direct traffic to the tunnel.

© 2018 Palo Alto Networks, Inc.

3. Enter the 64-bit extended unique **Interface ID** in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. By default, the firewall will use the EUI-64 generated from the physical interface's MAC address.
4. To assign an IPv6 **Address** to the tunnel interface, **Add** the IPv6 address and prefix length, for example 2001:400:f00::1/64. If Prefix is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.
    1. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address.
    2. Select **Anycast** to include routing through the nearest node.

STEP 5 | Set up key exchange.

Configure one of the following types of key exchange:

Set up Auto Key exchange

1. Select the IKE Gateway. To set up an IKE gateway, see Set Up an IKE Gateway.
2. (Optional) Select the default IPSec Crypto Profile. To create a new IPSec Profile, see Define IPSec Crypto Profiles.

Set up Manual Key exchange

1. Specify the **SPI** for the local firewall. SPI is a 32-bit hexadecimal index that is added to the header for IPSec tunneling to assist in differentiating between IPSec traffic flows; it is used to create the SA required for establishing a VPN tunnel.
2. Select the **Interface** that will be the tunnel endpoint, and optionally select the IP address for the local interface that is the endpoint of the tunnel.
3. Select the protocol to be used—**AH** or **ESP**.
4. For AH, select the **Authentication** method from the drop-down and enter a **Key** and then **Confirm Key**.
5. For ESP, select the **Authentication** method from the drop-down and enter a **Key** and then **Confirm Key**. Then, select the **Encryption** method and enter a **Key** and then **Confirm Key**, if needed.
6. Specify the **SPI** for the remote peer.
7. Enter the **Remote Address**, the IP address of the remote peer.

STEP 6 | Protect against a replay attack.

A replay attack occurs when a packet is maliciously intercepted and retransmitted by the interceptor.

Select the **Show Advanced Options** check box, select **Enable Replay Protection** to detect and neutralize against replay attacks.

STEP 7 | (Optional) Preserve the Type of Service header for the priority or treatment of IP packets.

In the Show Advanced Options section, select **Copy TOS Header**. This copies the Type of Service (TOS) header from the inner IP header to the outer IP header of the encapsulated packets in order to preserve the original TOS information.



If there are multiple sessions inside the tunnel (each with a different TOS value), copying the TOS header can cause the IPSec packets to arrive out of order.

STEP 8 | Enable Tunnel Monitoring.



You must assign an IP address to the tunnel interface for monitoring.

© 2018 Palo Alto Networks, Inc.

To alert the device administrator to tunnel failures and to provide automatic failover to another tunnel interface:

1. Specify a **Destination IP** address on the other side of the tunnel to determine if the tunnel is working properly.
2. Select a **Profile** to determine the action on tunnel failure. To create a new profile, see Define a Tunnel Monitoring Profile.

STEP 9 | Create a Proxy ID to identify the VPN peers.

This step is required only if the VPN peer uses policy-based VP).

1. Select **Network** > **IPSec Tunnels** and click **Add**.
2. Select the **Proxy IDs** tab.
3. Select the **IPv4** or **IPv6** tab.
4. Click **Add** and enter the **Proxy ID** name.
5. Enter the **Local** IP address or subnet for the VPN gateway.
6. Enter the **Remote** address for the VPN gateway.
7. Select the **Protocol** from the drop-down:
   - **Number**—Specify the protocol number (used for interoperability with third-party devices).
   - **Any**—Allows TCP and/or UDP traffic.
   - **TCP**—Specify the Local Port and Remote Port numbers.
   - **UDP**—Specify the Local Port and Remote Port numbers.
8. Click **OK**.

STEP 10 | Commit your changes.

Click **OK** and **Commit**.

Set Up Tunnel Monitoring

To provide uninterrupted VPN service, you can use the Dead Peer Detection capability along with the tunnel monitoring capability on the firewall. You can also monitor the status of the tunnel. These monitoring tasks are described in the following sections:

- Define a Tunnel Monitoring Profile on page 747
- View the Status of the Tunnels on page 748

Define a Tunnel Monitoring Profile

A tunnel monitoring profile allows you to verify connectivity between the VPN peers; you can configure the tunnel interface to ping a destination IP address at a specified interval and specify the action if the communication across the tunnel is broken.

STEP 1 | Select **Network** > **Network Profiles** > **Monitor**. A default tunnel monitoring profile is available for use.

STEP 2 | Click **Add**, and enter a **Name** for the profile.

STEP 3 | Select the **Action** to take if the destination IP address is unreachable.

- **Wait Recover**—the firewall waits for the tunnel to recover. It continues to use the tunnel interface in routing decisions as if the tunnel were still active.
- **Fail Over**—forces traffic to a back-up path if one is available. The firewall disables the tunnel interface, and thereby disables any routes in the routing table that use the interface.

© 2018 Palo Alto Networks, Inc.

In either case, the firewall attempts to accelerate the recovery by negotiating new IPSec keys.

STEP 4 | Specify the **Interval (sec)** and **Threshold** to trigger the specified action.

- **Threshold** specifies the number of heartbeats to wait before taking the specified action (range is 2-100; default is 5).
- **Interval (sec)** specifies the time (in seconds) between heartbeats (range is 2-10; default is 3).

STEP 5 | Attach the monitoring profile to the IPsec Tunnel configuration. See Enable Tunnel Monitoring.

View the Status of the Tunnels

The status of the tunnel informs you about whether or not valid IKE phase-1 and phase-2 SAs have been established, and whether the tunnel interface is up and available for passing traffic.

Because the tunnel interface is a logical interface, it cannot indicate a physical link status. Therefore, you must enable tunnel monitoring so that the tunnel interface can verify connectivity to an IP address and determine if the path is still usable. If the IP address is unreachable, the firewall will either wait for the tunnel to recover or failover. When a failover occurs, the existing tunnel is torn down and routing changes are triggered to set up a new tunnel and redirect traffic.

STEP 1 | Select **Network** > **IPSec Tunnels**.

STEP 2 | View the **Tunnel Status**.

- Green indicates a valid IPSec SA tunnel.
- Red indicates that IPSec SA is not available or has expired.

STEP 3 | View the **IKE Gateway Status**.

- Green indicates a valid IKE phase-1 SA.
- Red indicates that IKE phase-1 SA is not available or has expired.

STEP 4 | View the **Tunnel Interface Status**.

- Green indicates that the tunnel interface is up.
- Red indicates that the tunnel interface is down, because tunnel monitoring is enabled and the status is down.

To troubleshoot a VPN tunnel that is not yet up, see Interpret VPN Error Messages on page 751.

Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel

You can enable, disable, refresh or restart an IKE gateway or VPN tunnel to make troubleshooting easier.

- Enable or Disable an IKE Gateway or IPSec Tunnel on page 748
- Refresh and Restart Behaviors on page 749
- Refresh or Restart an IKE Gateway or IPSec Tunnel on page 749

Enable or Disable an IKE Gateway or IPSec Tunnel

Enable or disable an IKE gateway or IPSec tunnel to make troubleshooting easier.

- Enable or disable an IKE gateway.

© 2018 Palo Alto Networks, Inc.

1. Select **Network** > **Network Profiles** > **IKE Gateways** and select the gateway you want to enable or disable.
2. At the bottom of the screen, click **Enable** or **Disable**.

- Enable or disable an IPSec tunnel.
  1. Select **Network** > **IPSec Tunnels** and select the tunnel you want to enable or disable.
  2. At the bottom of the screen, click **Enable** or **Disable**.

Refresh and Restart Behaviors

You can Refresh or Restart an IKE Gateway or IPSec Tunnel. The refresh and restart behaviors for an IKE gateway and IPSec tunnel are as follows:

Phase	Refresh	Restart
IKE Gateway (IKE Phase 1)	Updates the onscreen statistics for the selected IKE gateway. Equivalent to issuing a second `show` command in the CLI (after an initial `show` command).	Restarts the selected IKE gateway. **IKEv2**: Also restarts any associated child IPSec security associations (SAs). **IKEv1**: Does not restart the associated IPSec SAs. A restart is disruptive to all existing sessions. Equivalent to issuing a **`clear`**, **`test`**, **`show`** command sequence in the CLI.
IPSec Tunnel (IKE Phase 2)	Updates the onscreen statistics for the selected IPSec tunnel. Equivalent to issuing a second `show` command in the CLI (after an initial `show` command).	Restarts the IPSec tunnel. A restart is disruptive to all existing sessions. Equivalent to issuing a **`clear`**, **`test`**, **`show`** command sequence in the CLI.

Refresh or Restart an IKE Gateway or IPSec Tunnel

Keep in mind that the result of restarting an IKE gateway depends on whether it is IKEv1 or IKEv2. See Refresh and Restart Behaviors for an IKE gateway (IKEv1 and IKEv2) and for an IPSec tunnel.

- Refresh or restart an IKE gateway.
  1. Select **Network** > **IPSec Tunnels** and select the tunnel for the gateway you want to refresh or restart.
  2. In the row for that tunnel, under the Status column, click **IKE Info**.
  3. At the bottom of the IKE Info screen, click the action you want:
     - **Refresh**—Updates the statistics on the screen.
     - **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is recreated.

- Refresh or restart an IPSec tunnel.

  You might determine that the tunnel needs to be refreshed or restarted because you use the tunnel monitor to monitor the tunnel status, or you use an external network monitor to monitor network connectivity through the IPSec tunnel.

© 2018 Palo Alto Networks, Inc.

1. Select **Network** > **IPSec Tunnels** and select the tunnel you want to refresh or restart.
2. In the row for that tunnel, under the Status column, click **Tunnel Info**.
3. At the bottom of the Tunnel Info screen, click the action you want:
   - **Refresh**—Updates the onscreen statistics.
   - **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is recreated.

Test VPN Connectivity

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 system log 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 system log 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
```


© 2018 Palo Alto Networks, Inc.

Interpret VPN Error Messages

The following table lists some of the common VPN error messages that are logged in the system log.

Table 7: Syslog Error Messages for VPN Issues

If error is this:	Try this:

`IKE phase-1 negotiation is failed as initiator, main mode. Failed SA: x.x.x.x[500]-y.y.y.y[500] cookie:84222f276c2fa2e9:0000000000000000 due to timeout.`
or
`IKE phase 1 negotiation is failed. Couldn't find configuration for IKE phase-1 request for peer IP x.x.x.x[1929]`
	• Verify that the public IP address for each VPN peer is accurate in the IKE Gateway configuration.
	• Verify that the IP addresses can be pinged and that routing issues are not causing the connection failure.

`Received unencrypted notify payload (no proposal chosen) from IP x.x.x.x[500] to y.y.y.y[500], ignored...`
or
`IKE phase-1 negotiation is failed. Unable to process peer's SA payload.`
	Check the IKE Crypto profile configuration to verify that the proposals on both sides have a common encryption, authentication, and DH Group proposal.

`pfs group mismatched:my: 2peer: 0`
or
`IKE phase-2 negotiation failed when processing SA payload. No suitable proposal found in peer's SA payload.`
	Check the IPSec Crypto profile configuration to verify that:
	• pfs is either enabled or disabled on both VPN peers
	• the DH Groups proposed by each peer has at least one DH Group in common

`IKE phase-2 negotiation failed when processing Proxy ID. Received local id x.x.x.x/x type IPv4 address protocol 0 port 0, received remote id y.y.y.y/y type IPv4 address protocol 0 port 0.`	The VPN peer on one end is using policy-based VPN. You must configure a Proxy ID on the Palo Alto Networks firewall. See Create a Proxy ID to identify the VPN peers..

© 2018 Palo Alto Networks, Inc.

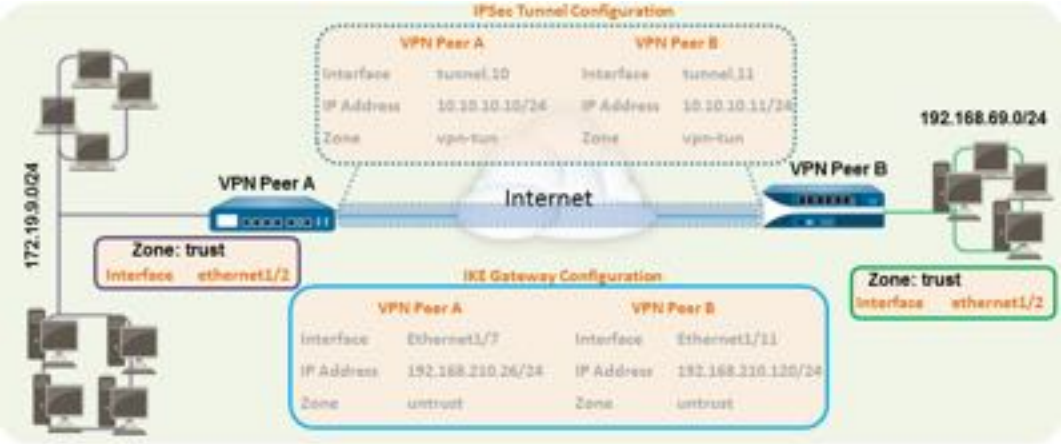
Site-to-Site VPN Quick Configs

The following sections provide instructions for configuring some common VPN deployments:

- Site-to-Site VPN with Static Routing on page 752
- Site-to-Site VPN with OSPF on page 755
- Site-to-Site VPN with Static and Dynamic Routing on page 758

Site-to-Site VPN with Static Routing

The following example shows a VPN connection between two sites that use static routes. Without dynamic routing, the tunnel interfaces on VPN Peer A and VPN Peer B do not require an IP address because the firewall automatically uses the tunnel interface as the next hop for routing traffic across the sites. However, to enable tunnel monitoring, a static IP address has been assigned to each tunnel interface.



STEP 1 | Configure a Layer 3 interface.

This interface is used for the IKE phase-1 tunnel.

1. Select **Network** > **Interfaces** > **Ethernet** and then select the interface you want to configure for VPN.
2. Select **Layer3** from the **Interface Type** drop-down.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
   - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
   - If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—192.168.210.26/24

   The configuration for VPN Peer B is:

© 2018 Palo Alto Networks, Inc.

- **Interface**—ethernet1/11
- **Security Zone**—untrust
- **Virtual Router**—default
- **IPv4**—192.168.210.120/24

STEP 2 | Create a tunnel interface and attach it to a virtual router and security zone.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.1**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.
4. Select the **Virtual Router**.
5. (Optional) Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface.

   With static routes, the tunnel interface does not require an IP address. For traffic that is destined to a specified subnet/IP address, the tunnel interface will automatically become the next hop. Consider adding an IP address if you want to enable tunnel monitoring.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—tunnel.10
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—172.19.9.2/24

   The configuration for VPN Peer B is:

   - **Interface**—tunnel.11
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—192.168.69.2/24

STEP 3 | Configure a static route, on the virtual router, to the destination subnet.

1. Select **Network** > **Virtual Router** and click the router you defined in the prior step.
2. Select **Static Route**, click **Add**, and enter a new route to access the subnet that is at the other end of the tunnel.

   In this example, the configuration for VPN Peer A is:

   - **Destination**—192.168.69.0/24
   - **Interface**—tunnel.10

   The configuration for VPN Peer B is:

   - **Destination**—172.19.9.0/24
   - **Interface**—tunnel.11

STEP 4 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).

Complete this task on both peers and make sure to set identical values.

1. Select **Network** > **Network Profiles** > **IKE Crypto**. In this example, we use the default profile.

© 2018 Palo Alto Networks, Inc.

2. Select **Network** > **Network Profiles** > **IPSec Crypto**. In this example, we use the default profile.

STEP 5 | Set up the IKE Gateway.

1. Select **Network** > **Network Profiles** > **IKE Gateway**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Local IP address**—192.168.210.26/24
   - **Peer IP type/address**—static/192.168.210.120
   - **Preshared keys**—enter a value
   - **Local identification**—None; this means that the local IP address will be used as the local identification value.
   - The configuration for VPN Peer B is:
   - **Interface**—ethernet1/11
   - **Local IP address**—192.168.210.120/24
   - **Peer IP type/address**—static/192.168.210.26
   - **Preshared keys**—enter same value as on Peer A
   - **Local identification**—None
3. Select **Advanced Phase 1 Options** and select the IKE Crypto profile you created earlier to use for IKE phase 1.

STEP 6 | Set up the IPSec Tunnel.

1. Select **Network** > **IPSec Tunnels**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Tunnel Interface**—tunnel.10
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IPSec Crypto profile defined in Step 4.

   The configuration for VPN Peer B is:

   - **Tunnel Interface**—tunnel.11
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IPSec Crypto defined in Step 4.
3. (Optional) Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity. Typically, the tunnel interface IP address for the VPN Peer is used.
4. (Optional) To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile.

STEP 7 | Create policies to allow traffic between the sites (subnets).

1. Select **Policies** > **Security**.
2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses.

STEP 8 | Commit any pending configuration changes.

Click **Commit**.

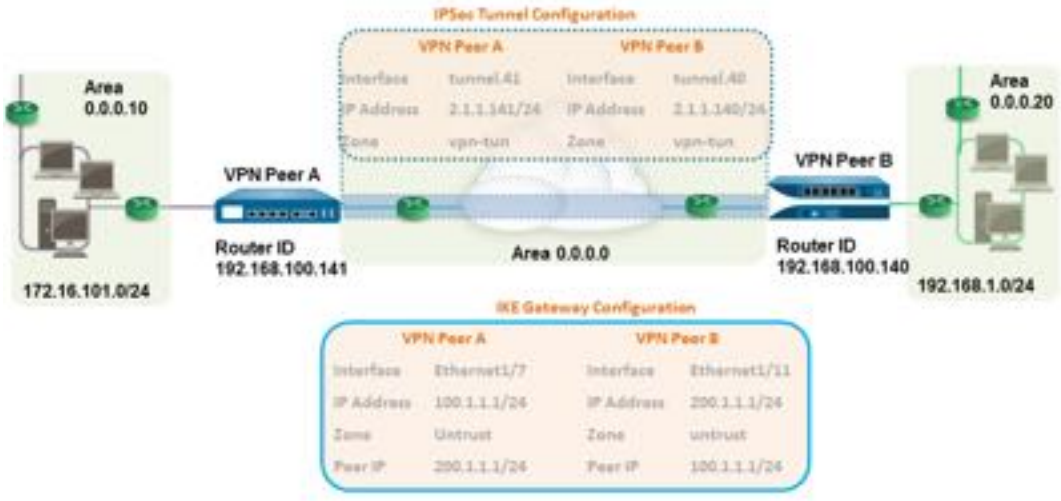
© 2018 Palo Alto Networks, Inc.

STEP 9 | Test VPN Connectivity.

See also View the Status of the Tunnels.

Site-to-Site VPN with OSPF

In this example, each site uses OSPF for dynamic routing of traffic. The tunnel IP address on each VPN peer is statically assigned and serves as the next hop for routing traffic between the two sites.



STEP 1 | Configure the Layer 3 interfaces on each firewall.

1. Select **Network** > **Interfaces** > **Ethernet** and then select the interface you want to configure for VPN.
2. Select **Layer3** from the **Interface Type** drop-down.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
   - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
   - If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—100.1.1.1/24

   The configuration for VPN Peer B is:

   - **Interface**—ethernet1/11
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—200.1.1.1/24

STEP 2 | Create a tunnel interface and attach it to a virtual router and security zone.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.

© 2018 Palo Alto Networks, Inc.

2. In the **Interface Name** field, specify a numeric suffix, such as, **.11**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example, vpn-tun), and then click **OK**.
4. Select the **Virtual Router**.
5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.

   This IP address will be used as the next hop IP address to route traffic to the tunnel and can also be used to monitor the status of the tunnel.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—tunnel.41
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—2.1.1.141/24

   The configuration for VPN Peer B is:

   - **Interface**—tunnel.40
   - **Security Zone**—vpn_tun
   - **Virtual Router**—default
   - **IPv4**—2.1.1.140/24

STEP 3 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).

Complete this task on both peers and make sure to set identical values.

1. Select **Network** > **Network Profiles** > **IKE Crypto**. In this example, we use the default profile.
2. Select **Network** > **Network Profiles** > **IPSec Crypto**. In this example, we use the default profile.

STEP 4 | Set up the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall.

For more information on the OSPF options that are available on the firewall, see Configure OSPF.

Use Broadcast as the link type when there are more than two OSPF routers that need to exchange routing information.

1. Select **Network** > **Virtual Routers**, and select the default router or add a new router.
2. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.
3. In this example, the OSPF configuration for VPN Peer A is:
   - **Router ID**: 192.168.100.141
   - **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p
   - **Area ID**: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast

   The OSPF configuration for VPN Peer B is:
   - **Router ID**: 192.168.100.140
   - **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p
   - **Area ID**: 0.0.0.20 that is assigned to the interface Ethernet1/15 and Link Type: Broadcast

© 2018 Palo Alto Networks, Inc.

STEP 5 | Set up the IKE Gateway.

This examples uses static IP addresses for both VPN peers. Typically, the corporate office uses a statically configured IP address, and the branch side can be a dynamic IP address; dynamic IP addresses are not best suited for configuring stable services such as VPN.

1. Select **Network** > **Network Profiles** > **IKE Gateway**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Local IP address**—100.1.1.1/24
   - **Peer IP address**—200.1.1.1/24
   - **Preshared keys**—enter a value

   The configuration for VPN Peer B is:

   - **Interface**—ethernet1/11
   - **Local IP address**—200.1.1.1/24
   - **Peer IP address**—100.1.1.1/24
   - **Preshared keys**—enter same value as on Peer A
3. Select the IKE Crypto profile you created earlier to use for IKE phase 1.

STEP 6 | Set up the IPSec Tunnel.

1. Select **Network** > **IPSec Tunnels**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Tunnel Interface**—tunnel.41
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IKE Gateway defined above.

   The configuration for VPN Peer B is:

   - **Tunnel Interface**—tunnel.40
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IKE Gateway defined above.
3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.
4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile.

STEP 7 | Create policies to allow traffic between the sites (subnets).

1. Select **Policies** > **Security**.
2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses.

STEP 8 | Verify OSPF adjacencies and routes from the CLI.

Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.

- **`show routing protocol ospf neighbor`**

© 2018 Palo Alto Networks, Inc.

```
admin@FW-A> show routing protocol ospf neighbor

  Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
           N/P:NSSA option, MC:multicase, E:AS external LSA capability, T:TOS capability
  ==========
  virtual router:               vr1
  neighbor address:             2.1.1.140
  local address binding:        0.0.0.0
  type:                         dynamic
  status:                       full
  neighbor router ID:           192.168.100.140
  area id:                      0.0.0.0
  neighbor priority:            1
  lifetime remain:              39
  messages pending:             0
  LSA request pending:          0
  options:                      0x42: O E
  hello suppressed:             no
```

```
admin@FW-B> show routing protocol ospf neighbor

  Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
           N/P:NSSA option, MC:multicase, E:AS external LSA capability, T:TOS capability
  ==========
  virtual router:               vr1
  neighbor address:             2.1.1.141
  local address binding:        0.0.0.0
  type:                         dynamic
  status:                       full
  neighbor router ID:           192.168.100.141
  area id:                      0.0.0.0
  neighbor priority:            1
  lifetime remain:              39
  messages pending:             0
  LSA request pending:          0
  options:                      0x42: O E
  hello suppressed:             no
```

- **show routing route type ospf**

```
admin@FW-A> show routing route type ospf

flags: A:active, 7:loose, C:connect, H:host, S:static, ~:internal, R:rip, O:ospf, B:bgp,
       Oi:ospf intra-area, Oo:ospf inter-area, O1:ospf ext-type-1, O2:ospf ext-type-2


VIRTUAL ROUTER: vr1 (id 1)
  ==========
destination             nexthop             metric flags      age   interface          next-AS
2.1.1.0/24              0.0.0.0                 10   Oi       6760  tunnel.41
172.16.101.0/24         0.0.0.0                 10   Oi       6854  ethernet1/1
192.168.1.0/24          2.1.1.140               20 A Oo       6754  tunnel.40
total routes shown: 3
```

```
admin@FW-B> show routing route type ospf

flags: A:active, C:connect, H:host, S:static, R:rip, O:ospf,
       Oi:ospf intra-area, Oo:ospf inter-area, O1:ospf ext-type-1, O2:ospf ext-type-2


VIRTUAL ROUTER: vr1 (id 1)
  ==========
destination             nexthop             metric flags      age   interface
2.1.1.0/24              0.0.0.0                 10   Oi       20033 tunnel.49
172.16.101.0/24         2.1.1.141               20   AOo      6896  tunnel.48
192.168.1.0/24          0.0.0.0                 10   Oi       8058  ethernet1/15
total routes shown: 3
```

STEP 9 | Test VPN Connectivity.

See Set Up Tunnel Monitoring and View the Status of the Tunnels.

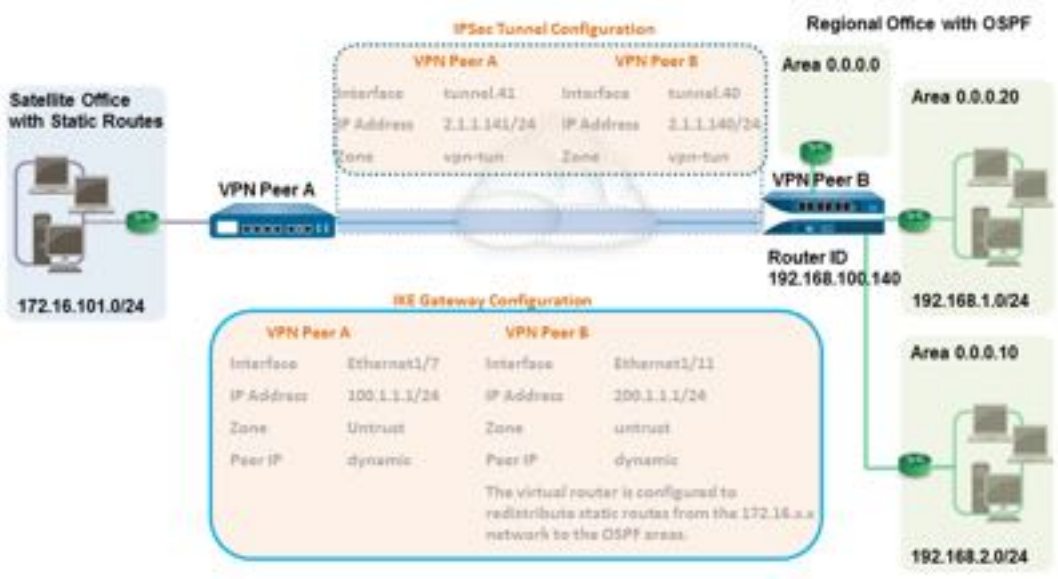
Site-to-Site VPN with Static and Dynamic Routing

In this example, one site uses static routes and the other site uses OSPF. When the routing protocol is not the same between the locations, the tunnel interface on each firewall must be configured with a static

© 2018 Palo Alto Networks, Inc.

IP address. Then, to allow the exchange of routing information, the firewall that participates in both the static and dynamic routing process 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 traffic destined to the 192.168.x.x network is routed to tunnel.41. The virtual router on VPN Peer B participates in both the static and the dynamic routing process and is configured with a redistribution profile in order to propagate (export) the static routes to the OSPF autonomous system.



STEP 1 | Configure the Layer 3 interfaces on each firewall.

1. Select **Network** > **Interfaces** > **Ethernet** and then select the interface you want to configure for VPN.
2. Select **Layer3** from the **Interface Type** drop-down.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
   - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
   - If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Security Zone**—untrust
   - **Virtual Router**—default
   - **IPv4**—100.1.1.1/24

   The configuration for VPN Peer B is:

   - **Interface**—ethernet1/11

© 2018 Palo Alto Networks, Inc.

- **Security Zone**—untrust
- **Virtual Router**—default
- **IPv4**—200.1.1.1/24

STEP 2 | Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).

Complete this task on both peers and make sure to set identical values.

1. Select **Network** > **Network Profiles** > **IKE Crypto**. In this example, we use the default profile.
2. Select **Network** > **Network Profiles** > **IPSec Crypto**. In this example, we use the default profile.

STEP 3 | Set up the IKE Gateway.

With pre-shared keys, to add authentication scrutiny when setting up the IKE phase-1 tunnel, you can set up Local and Peer Identification attributes and a corresponding value that is matched in the IKE negotiation process.

1. Select **Network** > **Network Profiles** > **IKE Gateway**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Interface**—ethernet1/7
   - **Local IP address**—100.1.1.1/24
   - **Peer IP type**—dynamic
   - **Preshared keys**—enter a value
   - **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer A.
   - **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer B

   The configuration for VPN Peer B is:

   - **Interface**—ethernet1/11
   - **Local IP address**—200.1.1.1/24
   - **Peer IP address**—dynamic
   - **Preshared keys**—enter same value as on Peer A
   - **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer B
   - **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer A
3. Select the IKE Crypto profile you created earlier to use for IKE phase 1.

STEP 4 | Create a tunnel interface and attach it to a virtual router and security zone.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, say, **.41**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.
4. Select the **Virtual Router**.
5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.

   This IP address will be used to route traffic to the tunnel and to monitor the status of the tunnel.
6. To save the interface configuration, click **OK**.

   In this example, the configuration for VPN Peer A is:

© 2018 Palo Alto Networks, Inc.

- **Interface**—tunnel.41
- **Security Zone**—vpn_tun
- **Virtual Router**—default
- **IPv4**—2.1.1.141/24

The configuration for VPN Peer B is:

- **Interface**—tunnel.42
- **Security Zone**—vpn_tun
- **Virtual Router**—default
- **IPv4**—2.1.1.140/24

STEP 5 | Specify the interface to route traffic to a destination on the 192.168.x.x network.

1. On VPN Peer A, select the virtual router.
2. Select **Static Routes**, and **Add** tunnel.41 as the **Interface** for routing traffic with a **Destination** in the 192.168.x.x network.

STEP 6 | Set up the static route and the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall.

1. On VPN Peer B, select **Network** > **Virtual Routers**, and select the default router or add a new router.
2. Select **Static Routes** and **Add** the tunnel IP address as the next hop for traffic in the 172.168.x.x. network.

   Assign the desired route metric; using a lower the value makes the a higher priority for route selection in the forwarding table.
3. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**.
4. In this example, the OSPF configuration for VPN Peer B is:

   - Router ID: 192.168.100.140
   - Area ID: 0.0.0.0 is assigned to the interface Ethernet 1/12 Link type: Broadcast
   - Area ID: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast
   - Area ID: 0.0.0.20 is assigned to the interface Ethernet1/15 and Link Type: Broadcast

STEP 7 | Create a redistribution profile to inject the static routes into the OSPF autonomous system.

1. Create a redistribution profile on VPN Peer B.
   1. Select **Network** > **Virtual Routers**, and select the router you used above.
   2. Select **Redistribution Profiles**, and click **Add**.
   3. Enter a Name for the profile and select **Redist** and assign a **Priority** value. If you have configured multiple profiles, the profile with the lowest priority value is matched first.
   4. Set **Source Type** as **static**, and click **OK**. The static route defined in Step 6-b will be used for the redistribution.
2. Inject the static routes in to the OSPF system.
   1. Select **OSPF** > **Export Rules** (for IPv4) or **OSPFv3** > **Export Rules** (for IPv6).
   2. Click **Add**, and select the redistribution profile that you just created.
   3. Select how the external routes are brought into the OSPF system. The default option, **Ext2** calculates the total cost of the route using only the external metrics. To use both internal and external OSPF metrics, use **Ext1**.
   4. Assign a **Metric** (cost value) for the routes injected into the OSPF system. This option allows you to change the metric for the injected route as it comes into the OSPF system.
   5. Click **OK**.

STEP 8 | Set up the IPSec Tunnel.

© 2018 Palo Alto Networks, Inc.

1. Select **Network** > **IPSec Tunnels**.
2. Click **Add** and configure the options in the **General** tab.

   In this example, the configuration for VPN Peer A is:

   - **Tunnel Interface**—tunnel.41
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IKE Gateway defined above.

   The configuration for VPN Peer B is:

   - **Tunnel Interface**—tunnel.40
   - **Type**—Auto Key
   - **IKE Gateway**—Select the IKE Gateway defined above.
   - **IPSec Crypto Profile**—Select the IKE Gateway defined above.
3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.
4. To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile.

STEP 9 | Create policies to allow traffic between the sites (subnets).

1. Select **Policies** > **Security**.
2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses.

STEP 10 | Verify OSPF adjacencies and routes from the CLI.

Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer.

- **`show routing protocol ospf neighbor`**

```
admin@FW-A> show routing protocol ospf neighbor

  Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
           N/P:NSSA option, MC:multicase, E:AS external LSA capability, T:TOS capability
  ==========
  virtual router:               vr1
  neighbor address:             2.1.1.140
  local address binding:        0.0.0.0
  type:                         dynamic
  status:                       full
  neighbor router ID:           192.168.100.140
  area id:                      0.0.0.0
  neighbor priority:            1
  lifetime remain:              39
  messages pending:             0
  LSA request pending:          0
  options:                      0x42: O E
  hello suppressed:             no
```


© 2018 Palo Alto Networks, Inc.

```
admin@FW-B> show routing protocol ospf neighbor

  Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
           N/P:NSSA option, MC:multicase, E:AS external LSA capability, T:TOS capability
  ==========
  virtual router:               vr1
  neighbor address:             2.1.1.141
  local address binding:        0.0.0.0
  type:                         dynamic
  status:                       full
  neighbor router ID:           192.168.100.141
  area id:                      0.0.0.0
  neighbor priority:            1
  lifetime remain:              39
  messages pending:             0
  LSA request pending:          0
  options:                      0x42: O E
  hello suppressed:             no
```

- **`show routing route`**

The following is an example of the output on each VPN peer.

```
VPN PeerA

destination          next hop        metric  flags   age     interface       next-AS
192.168.1.0/24       2.1.1.141       20      A S             tunnel.41
192.168.2.0/24       2.1.1.141       20      A S             tunnel.41
172.16.101.0/24      0.0.0.0         1       A H             ethernet1/1
2.1.1.140/24         2.1.1.141       20      A S             tunnel.41

VPN PeerB

destination          next hop        metric  flags   age     interface       next-AS
192.168.1.0/24       0.0.0.0         10      A Oo            ethernet1/1
192.168.2.0/24       0.0.0.0         10      A Oo            ethernet1/15
172.16.101.0/24      2.1.1.140       20      A H             tunnel.40
2.1.1.141/24         2.1.1.140       10      A C             tunnel.40
```

STEP 11 | Test VPN Connectivity.

See Set Up Tunnel Monitoring and View the Status of the Tunnels.

© 2018 Palo Alto Networks, Inc.

Large Scale VPN (LSVPN)

The GlobalProtect Large Scale VPN (LSVPN) feature on the Palo Alto Networks next-generation firewall simplifies the deployment of traditional hub and spoke VPNs, enabling you to quickly deploy enterprise networks with several branch offices with a minimum amount of configuration required on the remote *satellites*. This solution uses certificates for firewall authentication and IPSec to secure data.

LSVPN enables site-to-site VPNs between Palo Alto Networks firewalls. To set up a site-to-site VPN between a Palo Alto Networks firewall and another device, see VPNs.

The following topics describe the LSVPN components and how to set them up to enable site-to-site VPN services between Palo Alto Networks firewalls:

> LSVPN Overview
> Create Interfaces and Zones for the LSVPN
> Enable SSL Between GlobalProtect LSVPN Components
> Configure the Portal to Authenticate Satellites
> Configure GlobalProtect Gateways for LSVPN
> Configure the GlobalProtect Portal for LSVPN
> Prepare the Satellite to Join the LSVPN
> Verify the LSVPN Configuration
> LSVPN Quick Configs

© 2018 Palo Alto Networks, Inc.

LSVPN Overview

GlobalProtect provides a complete infrastructure for managing secure access to corporate resources from your remote sites. This infrastructure includes the following components:

- **GlobalProtect Portal**—Provides the management functions for your GlobalProtect LSVPN infrastructure. Every satellite that participates in the GlobalProtect LSVPN receives configuration information from the portal, including configuration information to enable the satellites (the spokes) to connect to the gateways (the hubs). You configure the portal on an interface on any Palo Alto Networks next-generation firewall.
- **GlobalProtect Gateways**—A Palo Alto Networks firewall that provides the tunnel end point for satellite connections. The resources that the satellites access is protected by security policy on the gateway. It is not required to have a separate portal and gateway; a single firewall can function both as portal and gateway.
- **GlobalProtect Satellite**—A Palo Alto Networks firewall at a remote site that establishes IPSec tunnels with the gateway(s) at your corporate office(s) for secure access to centralized resources. Configuration on the satellite firewall is minimal, enabling you to quickly and easily scale your VPN as you add new sites.

The following diagram illustrates how the GlobalProtect LSVPN components work together.



© 2018 Palo Alto Networks, Inc.

Create Interfaces and Zones for the LSVPN

You must configure the following interfaces and zones for your LSVPN infrastructure:

- **GlobalProtect portal**—Requires a Layer 3 interface for GlobalProtect satellites to connect to. If the portal and gateway are on the same firewall, they can use the same interface. The portal must be in a zone that is accessible from your branch offices.
- **GlobalProtect gateways**—Requires three interfaces: a Layer 3 interface in the zone that is reachable by the remote satellites, an internal interface in the trust zone that connects to the protected resources, and a logical tunnel interface for terminating the VPN tunnels from the satellites. Unlike other site-to-site VPN solutions, the GlobalProtect gateway only requires a single tunnel interface, which it will use for tunnel connections with all of your remote satellites (point-to-multi-point). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface. GlobalProtect supports both IPv6 and IPv4 addressing for the tunnel interface.
- **GlobalProtect satellites**—Requires a single tunnel interface for establishing a VPN with the remote gateways (up to a maximum of 25 gateways). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface. GlobalProtect supports both IPv6 and IPv4 addressing for the tunnel interface.

For more information about portals, gateways, and satellites see LSVPN Overview.

STEP 1 | Configure a Layer 3 interface.

The portal and each gateway and satellite all require a Layer 3 interface to enable traffic to be routed between sites.

If the gateway and portal are on the same firewall, you can use a single interface for both components.

1. Select **Network** > **Interfaces** > **Ethernet** and then select the interface you want to configure for GlobalProtect LSVPN.
2. Select **Layer3** from the **Interface Type** drop-down.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
   - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
   - If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. Assign an IP address to the interface:
   - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
   - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
6. To save the interface configuration, click **OK**.

STEP 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.



IP addresses are not required on the tunnel interface unless you plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.

© 2018 Palo Alto Networks, Inc.



Make sure to enable User-ID in the zone where the VPN tunnels terminate.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
   - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
   - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *lsvpn-tun*), select the **Enable User Identification** check box, and then click **OK**.
4. Select the **Virtual Router**.
5. (Optional) To assign an IP address to the tunnel interface:
   - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
   - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
6. To save the interface configuration, click **OK**.

STEP 3 | If you created a separate zone for tunnel termination of VPN connections, create a security policy to enable traffic flow between the VPN zone and your trust zone.

For example, a policy rule enables traffic between the *lsvpn-tun* zone and the *L3-Trust* zone.

STEP 4 | Commit your changes.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Enable SSL Between GlobalProtect LSVPN Components

All interaction between the GlobalProtect components occurs over an SSL/TLS connection. Therefore, you must generate and/or install the required certificates before configuring each component so that you can reference the appropriate certificate(s) and/or certificate profiles in the configurations for each component. The following sections describe the supported methods of certificate deployment, descriptions and best practice guidelines for the various GlobalProtect certificates, and provide instructions for generating and deploying the required certificates:

- About Certificate Deployment on page 770
- Deploy Server Certificates to the GlobalProtect LSVPN Components on page 770
- Deploy Client Certificates to the GlobalProtect Satellites Using SCEP on page 772

About Certificate Deployment

There are two basic approaches to deploying certificates for GlobalProtect LSVPN:

- **Enterprise Certificate Authority**—If you already have your own enterprise certificate authority, you can use this internal CA to issue an intermediate CA certificate for the GlobalProtect portal to enable it to issue certificates to the GlobalProtect gateways and satellites. You can also configure the GlobalProtect portal to act as a Simple Certificate Enrollment Protocol (SCEP) client to issue client certificates to GlobalProtect satellites.
- **Self-Signed Certificates**—You can generate a self-signed root CA certificate on the firewall and use it to issue server certificates for the portal, gateway(s), and satellite(s). As a best practice, create a self-signed root CA certificate on the portal and use it to issue server certificates for the gateways and satellites. This way, the private key used for certificate signing stays on the portal.

Deploy Server Certificates to the GlobalProtect LSVPN Components

The GlobalProtect LSVPN components use SSL/TLS to mutually authenticate. Before deploying the LSVPN, you must assign an SSL/TLS service profile to each portal and gateway. The profile specifies the server certificate and allowed TLS versions for communication with satellites. You don't need to create SSL/TLS service profiles for the satellites because the portal will issue a server certificate for each satellite during the first connection as part of the satellite registration process.

In addition, you must import the root certificate authority (CA) certificate used to issue the server certificates onto each firewall that you plan to host as a gateway or satellite. Finally, on each gateway and satellite participating in the LSVPN, you must configure a certificate profile that will enable them to establish an SSL/TLS connection using mutual authentication.

The following workflow shows the best practice steps for deploying SSL certificates to the GlobalProtect LSVPN components:

STEP 1 | On the firewall hosting the GlobalProtect portal, create the root CA certificate for signing the certificates of the GlobalProtect components.

Create a Self-Signed Root CA Certificate:

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Generate**.
2. Enter a **Certificate Name**, such as **`LSVPN_CA`**.

© 2018 Palo Alto Networks, Inc.

3. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed).
4. Select the **Certificate Authority** check box and then click **OK** to generate the certificate.

STEP 2 | Create SSL/TLS service profiles for the GlobalProtect portal and gateways.

For the portal and each gateway, you must assign an SSL/TLS service profile that references a unique self-signed server certificate.

The best practice is to issue all of the required certificates on the portal, so that the signing certificate (with the private key) doesn't have to be exported.

If the GlobalProtect portal and gateway are on the same firewall interface, you can use the same server certificate for both components.

1. Use the root CA on the portal to Generate a Certificate for each gateway you will deploy:
    1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Generate**.
    2. Enter a **Certificate Name**.
    3. Enter the FQDN (recommended) or IP address of the interface where you plan to configure the gateway in the **Common Name** field.
    4. In the **Signed By** field, select the **LSVPN_CA** certificate you just created.
    5. In the Certificate Attributes section, click **Add** and define the attributes to uniquely identify the gateway. If you add a **Host Name** attribute (which populates the SAN field of the certificate), it must exactly match the value you defined for the **Common Name**.
    6. **Generate** the certificate.
2. Configure an SSL/TLS Service Profile for the portal and each gateway:
    1. Select **Device** > **Certificate Management** > **SSL/TLS Service Profile** and click **Add**.
    2. Enter a **Name** to identify the profile and select the server **Certificate** you just created for the portal or gateway.
    3. Define the range of TLS versions (**Min Version** to **Max Version**) allowed for communicating with satellites and click **OK**.

STEP 3 | Deploy the self-signed server certificates to the gateways.

Best Practices:

- Export the self-signed server certificates issued by the root CA from the portal and import them onto the gateways.
- Be sure to issue a unique server certificate for each gateway.
- The Common Name (CN) and, if applicable, the Subject Alternative Name (SAN) fields of the certificate must match the IP address or fully qualified domain name (FQDN) of the interface where you configure the gateway.

1. On the portal, select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**, select the gateway certificate you want to deploy, and click **Export**.
2. Select **Encrypted Private Key and Certificate (PKCS12)** from the **File Format** drop-down.
3. Enter (and re-enter) a **Passphrase** to encrypt the private key associated with the certificate and then click **OK** to download the PKCS12 file to your computer.
4. On the gateway, select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
5. Enter a **Certificate Name**.

© 2018 Palo Alto Networks, Inc.

6. Enter the path and name to the **Certificate File** you just downloaded from the portal, or **Browse** to find the file.
7. Select **Encrypted Private Key and Certificate (PKCS12)** as the **File Format**.
8. Enter the path and name to the PKCS12 file in the **Key File** field or **Browse** to find it.
9. Enter and re-enter the **Passphrase** you used to encrypt the private key when you exported it from the portal and then click **OK** to import the certificate and key.

STEP 4 | Import the root CA certificate used to issue server certificates for the LSVPN components.

You must import the root CA certificate onto all gateways and satellites. For security reasons, make sure you export the certificate only, and not the associated private key.

1. Download the root CA certificate from the portal.
    1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
    2. Select the root CA certificate used to issue certificates for the LSVPN components and click **Export**.
    3. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (Do not export the private key.)
2. On the firewalls hosting the gateways and satellites, import the root CA certificate.
    1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
    2. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.
    3. **Browse** to the **Certificate File** you downloaded from the CA.
    4. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.
    5. Select the certificate you just imported on the **Device Certificates** tab to open it.
    6. Select **Trusted Root CA** and then click **OK**.
    7. **Commit** the changes.

STEP 5 | Create a certificate profile.

The GlobalProtect LSVPN portal and each gateway require a certificate profile that specifies which certificate to use to authenticate the satellites.

1. Select **Device** > **Certificate Management** > **Certificate Profile** and click **Add** and enter a profile **Name**.
2. Make sure **Username Field** is set to **None**.
3. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you imported in the previous step.
4. (Recommended) Enable use of CRL and/or OCSP to enable certificate status verification.
5. Click **OK** to save the profile.

STEP 6 | Commit your changes.

Click **Commit**.

Deploy Client Certificates to the GlobalProtect Satellites Using SCEP

As an alternative method for deploying client certificates to satellites, you can configure your GlobalProtect portal to act as a Simple Certificate Enrollment Protocol (SCEP) client to a SCEP server in your enterprise PKI. SCEP operation is dynamic in that the enterprise PKI generates a certificate when the portal requests it and sends the certificate to the portal.

© 2018 Palo Alto Networks, Inc.

When the satellite device requests a connection to the portal or gateway, it also includes its serial number with the connection request. The portal submits a CSR to the SCEP server using the settings in the SCEP profile and automatically includes the serial number of the device in the subject of the client certificate. After receiving the client certificate from the enterprise PKI, the portal transparently deploys the client certificate to the satellite device. The satellite device then presents the client certificate to the portal or gateway for authentication.

STEP 1 | Create a SCEP profile.

1. Select **Device** > **Certificate Management** > **SCEP** and then **Add** a new profile.
2. Enter a **Name** to identify the SCEP profile.
3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the **Location** where the profile is available.

STEP 2 | (Optional) To make the SCEP-based certificate generation more secure, configure a SCEP challenge-response mechanism between the PKI and portal for each certificate request.

After you configure this mechanism, its operation is invisible, and no further input from you is necessary.

To comply with the U.S. Federal Information Processing Standard (FIPS), use a **Dynamic** SCEP challenge and specify a **Server URL** that uses HTTPS (see Step 7).

Select one of the following options:

- **None**—(Default) The SCEP server does not challenge the portal before it issues a certificate.
- **Fixed**—Obtain the enrollment challenge password from the SCEP server (for example, **`http://10.200.101.1/CertSrv/mscep_admin/`**) in the PKI infrastructure and then copy or enter the password into the Password field.
- **Dynamic**—Enter the SCEP **Server URL** where the portal-client submits these credentials (for example, **`http://10.200.101.1/CertSrv/mscep_admin/`**), and a username and OTP of your choice. The username and password can be the credentials of the PKI administrator.

STEP 3 | Specify the settings for the connection between the SCEP server and the portal to enable the portal to request and receive client certificates.

To identify the satellite, the portal automatically includes the device serial number in the CSR request to the SCEP server. Because the SCEP profile requires a value in the **Subject** field, you can leave the default **`$USERNAME`** token even though the value is not used in client certificates for LSVPN.

1. Configure the **Server URL** that the portal uses to reach the SCEP server in the PKI (for example, **`http://10.200.101.1/certsrv/mscep/`**).
2. Enter a string (up to 255 characters in length) in the **CA-IDENT Name** field to identify the SCEP server.
3. Select the **Subject Alternative Name Type**:
   - **RFC 822 Name**—Enter the email name in a certificate's subject or Subject Alternative Name extension.
   - **DNS Name**—Enter the DNS name used to evaluate certificates.
   - **Uniform Resource Identifier**—Enter the name of the resource from which the client will obtain the certificate.
   - **None**—Do not specify attributes for the certificate.

STEP 4 | (Optional) Configure cryptographic settings for the certificate.

- Select the key length (**Number of Bits**) for the certificate. If the firewall is in FIPS-CC mode and the key generation algorithm is RSA. The RSA keys must be 2048 bits or larger.
- Select the **Digest for CSR** which indicates the digest algorithm for the certificate signing request (CSR): SHA1, SHA256, SHA384, or SHA512.

© 2018 Palo Alto Networks, Inc.

STEP 5 | (Optional) Configure the permitted uses of the certificate, either for signing or encryption.

- To use this certificate for signing, select the **Use as digital signature** check box. This enables the endpoint use the private key in the certificate to validate a digital signature.
- To use this certificate for encryption, select the **Use for key encipherment** check box. This enables the client use the private key in the certificate to encrypt data exchanged over the HTTPS connection established with the certificates issued by the SCEP server.

STEP 6 | (Optional) To ensure that the portal is connecting to the correct SCEP server, enter the **CA Certificate Fingerprint**. Obtain this fingerprint from the SCEP server interface in the Thumbprint field.

1. Enter the URL for the SCEP server's administrative UI (for example, **`http://<hostname or IP>/CertSrv/mscep_admin/`**).
2. Copy the thumbprint and enter it in the **CA Certificate Fingerprint** field.

STEP 7 | Enable mutual SSL authentication between the SCEP server and the GlobalProtect portal. This is required to comply with the U.S. Federal Information Processing Standard (FIPS).



FIPS-CC operation is indicated on the firewall login page and in its status bar.

Select the SCEP server's root **CA Certificate**. Optionally, you can enable mutual SSL authentication between the SCEP server and the GlobalProtect portal by selecting a **Client Certificate**.

STEP 8 | Save and commit the configuration.

1. Click **OK** to save the settings and close the SCEP configuration.
2. **Commit** the configuration.

The portal attempts to request a CA certificate using the settings in the SCEP profile and saves it to the firewall hosting the portal. If successful, the CA certificate is shown in **Device** > **Certificate Management** > **Certificates**.

STEP 9 | (Optional) If after saving the SCEP profile, the portal fails to obtain the certificate, you can manually generate a certificate signing request (CSR) from the portal.

1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and then click **Generate**.
2. Enter a **Certificate Name**. This name cannot contain spaces.
3. Select the **SCEP Profile** to use to submit a CSR to your enterprise PKI.
4. Click **OK** to submit the request and generate the certificate.

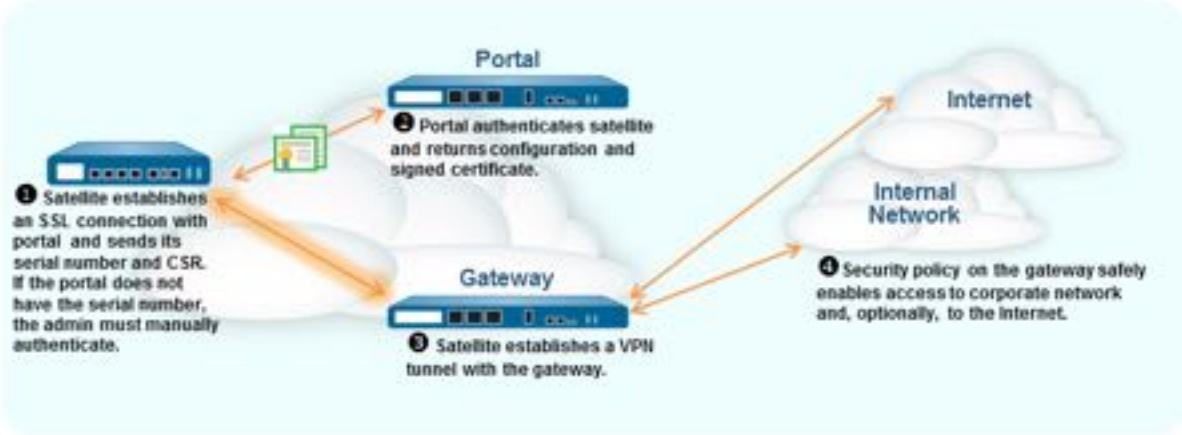
© 2018 Palo Alto Networks, Inc.

Configure the Portal to Authenticate Satellites

In order to register with the LSVPN, each satellite must establish an SSL/TLS connection with the portal. After establishing the connection, the portal authenticates the satellite to ensure that is authorized to join the LSVPN. After successfully authenticating the satellite, the portal will issue a server certificate for the satellite and push the LSVPN configuration specifying the gateways to which the satellite can connect and the root CA certificate required to establish an SSL connection with the gateways.

There are two ways that the satellite can authenticate to the portal during its initial connection:

- **Serial number**—You can configure the portal with the serial number of the satellite firewalls that are authorized to join the LSVPN. During the initial satellite connection to the portal, the satellite presents its serial number to the portal and if the portal has the serial number in its configuration, the satellite will be successfully authenticated. You add the serial numbers of authorized satellites when you configure the portal. See Configure the Portal.
- **Username and password**—If you would rather provision your satellites without manually entering the serial numbers of the satellites into the portal configuration, you can instead require the satellite administrator to authenticate when establishing the initial connection to the portal. Although the portal will always look for the serial number in the initial request from the satellite, if it cannot identify the serial number, the satellite administrator must provide a username and password to authenticate to the portal. Because the portal will always fall back to this form of authentication, you must create an authentication profile in order to commit the portal configuration. This requires that you set up an authentication profile for the portal LSVPN configuration even if you plan to authenticate satellites using the serial number.



The following workflow describes how to set up the portal to authenticate satellites against an existing authentication service. GlobalProtect LSVPN supports external authentication using a local database, LDAP (including Active Directory), Kerberos, TACACS+, or RADIUS.

STEP 1 | (External authentication only) Create a server profile on the portal.

The server profile defines how the firewall connects to an external authentication service to validate the authentication credentials that the satellite administrator enters.

If you use local authentication, skip this step and instead add a local user for the satellite administrator: see Add the user account to the local database.

Configure a server profile for the authentication service type:

© 2018 Palo Alto Networks, Inc.

- Add a RADIUS server profile.



You can use RADIUS to integrate with a Multi-Factor Authentication *service.*

- Add a TACACS+ server profile.
- Add a SAML IdP server profile.
- Add a Kerberos server profile.
- Add an LDAP server profile. If you use LDAP to connect to Active Directory (AD), create a separate LDAP server profile for every AD domain.

STEP 2 | Configure an authentication profile.

The authentication profile defines which server profile to use to authenticate satellites.

1. Select **Device** > **Authentication Profile** and click **Add**.
2. Enter a **Name** for the profile and then select the authentication **Type**. If the **Type** is an external service, select the **Server Profile** you created in the previous step. If you added a local user instead, set the **Type** to **Local Database**.
3. Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Configure GlobalProtect Gateways for LSVPN

Because the GlobalProtect configuration that the portal delivers to the satellites includes the list of gateways the satellite can connect to, it is a good idea to configure the gateways before configuring the portal.

Before you can configure the GlobalProtect gateway, you must complete the following tasks:

- Create Interfaces and Zones for the LSVPN on page 768 on the interface where you will configure each gateway. You must configure both the physical interface and the virtual tunnel interface.
- Enable SSL Between GlobalProtect LSVPN Components on page 770 by configuring the gateway server certificates, SSL/TLS service profiles, and certificate profile required to establish a mutual SSL/TLS connection from the GlobalProtect satellites to the gateway.

Configure each GlobalProtect gateway to participate in the LSVPN as follows:

STEP 1 | Add a gateway.

1. Select **Network** > **GlobalProtect** > **Gateways** and click **Add**.
2. In the **General** screen, enter a **Name** for the gateway. The gateway name should have no spaces and, as a best practice, should include the location or other descriptive information to help users and administrators identify the gateway.
3. (Optional) Select the virtual system to which this gateway belongs from the **Location** field.

STEP 2 | Specify the network information that enables satellite devices to connect to the gateway.

If you haven't created the network interface for the gateway, see Create Interfaces and Zones for the LSVPN on page 768 for instructions.

1. Select the **Interface** that satellites will use for ingress access to the gateway.
2. Specify the **IP Address Type** and **IP address** for gateway access:
   - The IP address type can be **IPv4** (only), **IPv6** (only), or **IPv4 and IPv6**. Use **IPv4 and IPv6** if your network supports dual stack configurations, where IPv4 and IPv6 run at the same time.
   - The IP address must be compatible with the IP address type. For example, **`172.16.1/0`** for IPv4 addresses or **`21DA:D3:0:2F3B`** for IPv6 addresses. For dual stack configurations, enter both an IPv4 and IPv6 address.
3. Click **OK** to save changes.

STEP 3 | Specify how the gateway authenticates satellites attempting to establish tunnels. If you haven't yet created an SSL/TLS Service profile for the gateway, see Deploy Server Certificates to the GlobalProtect LSVPN Components on page 770.

If you haven't set up the authentication profiles or certificate profiles, see Configure the Portal to Authenticate Satellites on page 775 for instructions.

If you have not yet set up the certificate profile, see Enable SSL Between GlobalProtect LSVPN Components on page 770 for instructions.

On the GlobalProtect Gateway Configuration dialog, select Authentication and then configure any of the following:

- To secure communication between the gateway and the satellites, select the **SSL/TLS Service Profile** for the gateway.
- To specify the authentication profile to use to authenticate satellites, **Add** a Client Authentication. Then, enter a **Name** to identify the configuration, select **OS**: **Satellite** to apply the configuration to all satellites, and specify the **Authentication Profile** to use to authenticate the satellite. You can also

© 2018 Palo Alto Networks, Inc.

select a **Certificate Profile** for the gateway to use to authenticate satellite devices attempting to establish tunnels.

STEP 4 | Configure the tunnel parameters and enable tunneling.

1. On the GlobalProtect Gateway Configuration dialog, select **Satellite** > **Tunnel Settings**.
2. Select the **Tunnel Configuration** check box to enable tunneling.
3. Select the **Tunnel Interface** you defined to terminate VPN tunnels established by the GlobalProtect satellites when you performed the task to Create Interfaces and Zones for the LSVPN on page 768.
4. (Optional) If you want to preserve the Type of Service (ToS) information in the encapsulated packets, select **Copy TOS**.



If there are multiple sessions inside the tunnel (each with a different TOS value), copying the TOS header can cause the IPSec packets to arrive out of order.

STEP 5 | (Optional) Enable tunnel monitoring.

Tunnel monitoring enables satellites to monitor its gateway tunnel connection, allowing it to failover to a backup gateway if the connection fails. Failover to another gateway is the only type of tunnel monitoring profile supported with LSVPN.

1. Select the **Tunnel Monitoring** check box.
2. Specify the **Destination IP Address** the satellites should use to determine if the gateway is active. You can specify an **IPv4** address, and **IPv6** address, or both. Alternatively, if you configured an IP address for the tunnel interface, you can leave this field blank and the tunnel monitor will instead use the tunnel interface to determine if the connection is active.
3. Select **Failover** from the **Tunnel Monitor Profile** drop-down (this is the only supported tunnel monitor profile for LSVPN).

STEP 6 | Select the IPSec Crypto profile to use when establishing tunnel connections.

The profile specifies the type of IPSec encryption and the authentication method for securing the data that will traverse the tunnel. Because both tunnel endpoints in an LSVPN are trusted firewalls within your organization, you can typically use the default (predefined) profile, which uses ESP as the IPSec protocol, group2 for the DH group, AES-128-CBC for encryption, and SHA-1 for authentication.

In the **IPSec Crypto Profile** drop-down, select **default** to use the predefined profile or select **New IPSec Crypto Profile** to define a new profile. For details on the authentication and encryption options, see Define IPSec Crypto Profiles on page 744.

STEP 7 | Configure the network settings to assign the satellites during establishment of the IPSec tunnel.



You can also configure the satellite to push the DNS settings to its local clients by configuring a DHCP server on the firewall hosting the satellite. In this configuration, the satellite will push DNS settings it learns from the gateway to the DHCP clients.

1. On the GlobalProtect Gateway Configuration dialog, select **Satellite** > **Network Settings**.
2. (Optional) If clients local to the satellite need to resolve FQDNs on the corporate network, configure the gateway to push DNS settings to the satellites in one of the following ways:
   - If the gateway has an interface that is configured as a DHCP client, you can set the **Inheritance Source** to that interface and assign the same settings received by the DHCP client to GlobalProtect satellites. You can also inherit the DNS suffix from the same source.
   - Manually define the **Primary DNS**, **Secondary DNS**, and **DNS Suffix** settings to push to the satellites.

© 2018 Palo Alto Networks, Inc.

3. To specify the **IP Pool** of addresses to assign the tunnel interface on the satellites when the VPN is established, click **Add** and then specify the IP address range(s) to use.
4. To define what destination subnets to route through the tunnel click **Add** in the **Access Route** area and then enter the routes as follows:
   - If you want to route all traffic from the satellites through the tunnel, leave this field blank.

     In this case, all traffic except traffic destined for the local subnet will be tunneled to the gateway.

   - To route only some traffic through the gateway (called *split tunneling*), specify the destination subnets that must be tunneled. In this case, the satellite will route traffic that is not destined for a specified access route using its own routing table. For example, you may choose to only tunnel traffic destined for your corporate network, and use the local satellite to safely enable Internet access.
   - If you want to enable routing between satellites, enter the summary route for the network protected by each satellite.

STEP 8 | (Optional) Define what routes, if any, the gateway will accept from satellites.

By default, the gateway will not add any routes satellites advertise to its routing table. If you do not want the gateway to accept routes from satellites, you do not need to complete this step.

1. To enable the gateway to accept routes advertised by satellites, select **Satellite** > **Route Filter**.
2. Select the **Accept published routes** check box.
3. To filter which of the routes advertised by the satellites to add to the gateway routing table, click **Add** and then define the subnets to include. For example, if all the satellites are configured with subnet 192.168.x.0/24 on the LAN side, configuring a permitted route of 192.168.0.0/16 to enable the gateway to only accept routes from the satellite if it is in the 192.168.0.0/16 subnet.

STEP 9 | Save the gateway configuration.

1. Click **OK** to save the settings and close the GlobalProtect Gateway Configuration dialog.
2. **Commit** the configuration.

© 2018 Palo Alto Networks, Inc.

Configure the GlobalProtect Portal for LSVPN

The GlobalProtect portal provides the management functions for your GlobalProtect LSVPN. Every satellite system that participates in the LSVPN receives configuration information from the portal, including information about available gateways as well as the certificate it needs in order to connect to the gateways.

The following sections provide procedures for setting up the portal:

- GlobalProtect Portal for LSVPN Prerequisite Tasks on page 780
- Configure the Portal on page 780
- Define the Satellite Configurations on page 781

GlobalProtect Portal for LSVPN Prerequisite Tasks

Before configuring the GlobalProtect portal, you must complete the following tasks:

- ❑ Create Interfaces and Zones for the LSVPN on page 768 on the interface where you will configure the portal.
- ❑ Enable SSL Between GlobalProtect LSVPN Components on page 770 by creating an SSL/TLS service profile for the portal server certificate, issuing gateway server certificates, and configuring the portal to issue server certificates for the GlobalProtect satellites.
- ❑ Configure the Portal to Authenticate Satellites on page 775 by defining the authentication profile that the portal will use to authenticate satellites if the serial number is not available.
- ❑ Configure GlobalProtect Gateways for LSVPN on page 777.

Configure the Portal

After you have completed the GlobalProtect Portal for LSVPN Prerequisite Tasks, configure the GlobalProtect portal as follows:

STEP 1 | Add the portal.

1. Select **Network** > **GlobalProtect** > **Portals** and click **Add**.
2. On the **General** tab, enter a **Name** for the portal. The portal name should not contain any spaces.
3. (Optional) Select the virtual system to which this portal belongs from the **Location** field.

STEP 2 | Specify the network information to enable satellites to connect to the portal.

If you haven't yet created the network interface for the portal, see Create Interfaces and Zones for the LSVPN for instructions.

1. Select the **Interface** that satellites will use for ingress access to the portal.
2. Specify the **IP Address Type** and **IP address** for satellite access to the portal:
   - The IP address type can be **IPv4** (for IPv4 traffic only), **IPv6** (for IPv6 traffic only, or **IPv4 and IPv6**. Use **IPv4 and IPv6** if your network supports dual stack configurations, where IPv4 and IPv6 run at the same time.
   - The IP address must be compatible with the IP address type. For example, **`172.16.1/0`** for IPv4 addresses or **`21DA:D3:0:2F3B`** for IPv6 addresses. For dual stack configurations, enter both an IPv4 and IPv6 address.
3. Click **OK** to save changes.

STEP 3 | Specify an SSL/TLS Service profile to use to enable the satellite to establish an SSL/TLS connection to the portal.

© 2018 Palo Alto Networks, Inc.

If you haven't yet created an SSL/TLS service profile for the portal and issued gateway certificates, see Deploy Server Certificates to the GlobalProtect LSVPN Components.

1. On the GlobalProtect Portal Configuration dialog, select **Authentication**.
2. Select the **SSL/TLS Service Profile**.

STEP 4 | Specify an authentication profile and optional certificate profile for authenticating satellites.

If the portal can't validate the serial numbers of connecting satellites, it will fall back to the authentication profile. Therefore, before you can save the portal configuration (by clicking OK), you must Configure an authentication profile.

Add a Client Authentication, and then enter a **Name** to identify the configuration, select **OS**: **Satellite** to apply the configuration to all satellites, and specify the **Authentication Profile** to use to authenticate satellite devices. You can also specify a **Certificate Profile** for the portal to use to authenticate satellite devices.

STEP 5 | Continue with defining the configurations to push to the satellites or, if you have already created the satellite configurations, save the portal configuration.

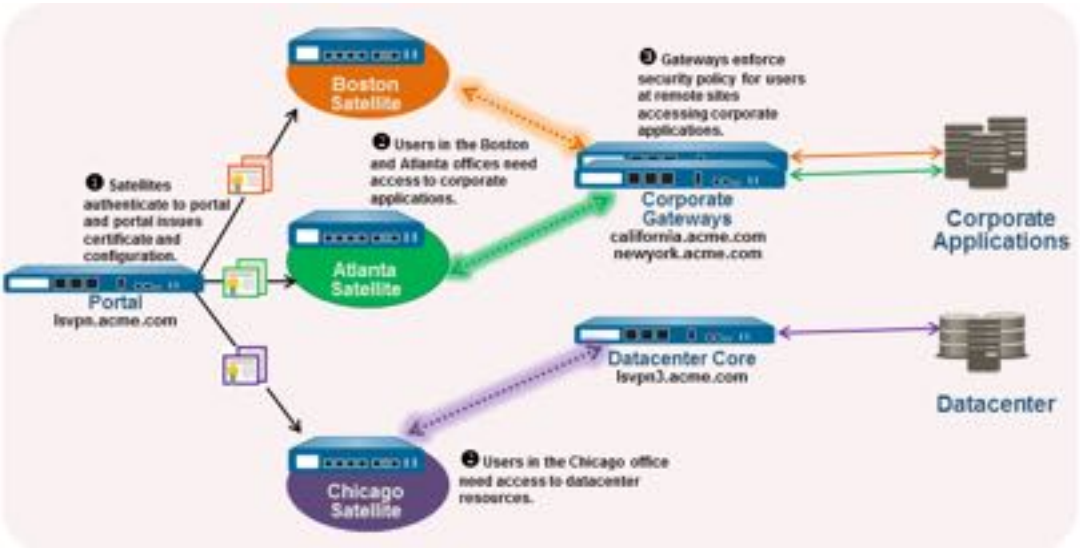
Click **OK** to save the portal configuration or continue to Define the Satellite Configurations.

Define the Satellite Configurations

When a GlobalProtect satellite connects and successfully authenticates to the GlobalProtect portal, the portal delivers a satellite configuration, which specifies what gateways the satellite can connect to. If all your satellites will use the same gateway and certificate configurations, you can create a single satellite configuration to deliver to all satellites upon successful authentication. However, if you require different satellite configurations—for example if you want one group of satellites to connect to one gateway and another group of satellites to connect to a different gateway—you can create a separate satellite configuration for each. The portal will then use the enrollment username/group name or the serial number of the satellite to determine which satellite configuration to deploy. As with security rule evaluation, the portal looks for a match starting from the top of the list. When it finds a match, it delivers the corresponding configuration to the satellite.

For example, the following figure shows a network in which some branch offices require VPN access to the corporate applications protected by your perimeter firewalls and another site needs VPN access to the data center.

© 2018 Palo Alto Networks, Inc.



Use the following procedure to create one or more satellite configurations.

STEP 1 | Add a satellite configuration.

The satellite configuration specifies the GlobalProtect LSVPN configuration settings to deploy to the connecting satellites. You must define at least one satellite configuration.

1. Select **Network** > **GlobalProtect** > **Portals** and select the portal configuration for which you want to add a satellite configuration and then select the **Satellite** tab.
2. In the Satellite section, click **Add**.
3. Enter a **Name** for the configuration.

   If you plan to create multiple configurations, make sure the name you define for each is descriptive enough to allow you to distinguish them.
4. To change how often a satellite should check the portal for configuration updates specify a value in the **Configuration Refresh Interval (hours)** field (range is 1-48; default is 24).

STEP 2 | Specify the satellites to which to deploy this configuration.

The portal uses the **Enrollment User/User Group** settings and/or **Devices** serial numbers to match a satellite to a configuration. Therefore, if you have multiple configurations, be sure to order them properly. As soon as the portal finds a match, it will deliver the configuration. Therefore, more specific configurations must precede more general ones. See Step 5 for instructions on ordering the list of satellite configurations.

Specify the match criteria for the satellite configuration as follows:

- To restrict this configuration to satellites with specific serial numbers, select the **Devices** tab, click **Add**, and enter serial number (you do not need to enter the satellite hostname; it will be automatically added when the satellite connects). Repeat this step for each satellite you want to receive this configuration.
- Select the **Enrollment User/User Group** tab, click **Add**, and then select the user or group you want to receive this configuration. Satellites that do not match on serial number will be required to authenticate as a user specified here (either an individual user or group member).

Before you can restrict the configuration to specific groups, you must Map Users to Groups.

STEP 3 | Specify the gateways that satellites with this configuration can establish VPN tunnels with.

© 2018 Palo Alto Networks, Inc.

Routes published by the gateway are installed on the satellite as static routes. The metric for the static route is 10x the routing priority. If you have more than one gateway, make sure to also set the routing priority to ensure that routes advertised by backup gateways have higher metrics compared to the same routes advertised by primary gateways. For example, if you set the routing priority for the primary gateway and backup gateway to 1 and 10 respectively, the satellite will use 10 as the metric for the primary gateway and 100 as the metric for the backup gateway.

1. On the **Gateways** tab, click **Add**.
2. Enter a descriptive **Name** for the gateway. The name you enter here should match the name you defined when you configured the gateway and should be descriptive enough identify the location of the gateway.
3. Enter the FQDN or IP address of the interface where the gateway is configured in the **Gateways** field. The address you specify must exactly match the Common Name (CN) in the gateway server certificate.
4. (Optional) If you are adding two or more gateways to the configuration, the **Routing Priority** helps the satellite pick the preferred gateway. Enter a value in the range of 1-25, with lower numbers having the higher priority (that is, the gateway the satellite will connect to if all gateways are available). The satellite will multiply the routing priority by 10 to determine the routing metric.

STEP 4 | Save the satellite configuration.

1. Click **OK** to save the satellite configuration.
2. If you want to add another satellite configuration, repeat the previous steps.

STEP 5 | Arrange the satellite configurations so that the proper configuration is deployed to each satellite.

- To move a satellite configuration up on the list of configurations, select the configuration and click **Move Up**.
- To move a satellite configuration down on the list of configurations, select the configuration and click **Move Down**.

STEP 6 | Specify the certificates required to enable satellites to participate in the LSVPN.

1. In the **Trusted Root CA** field, click **Add** and then select the CA certificate used to issue the gateway server certificates. The portal will deploy the root CA certificate you add here to all satellites as part of the configuration to enable the satellite to establish an SSL connection with the gateways. As a best practice, all of your gateways should use the same issuer.
2. Select the method of **Client Certificate** distribution:
   - **To store the client certificates on the portal**—select **Local** and select the Root CA certificate that the portal will use to issue client certificates to satellites upon successfully authenticating them from the **Issuing Certificate** drop-down.

     If the root CA certificate used to issue your gateway server certificates is not on the portal, you can Import it now. See Enable SSL Between GlobalProtect LSVPN Components *for details on how to import a root CA certificate.*

   - **To enable the portal to act as a SCEP client to dynamically request and issue client certificates**—select **SCEP** and then select the **SCEP** profile used to generate CSRs to your SCEP server.

     If the you have not yet set up the portal to act as a SCEP client, you can add a New SCEP profile now. See Deploy Client Certificates to the GlobalProtect Satellites Using SCEP *for details.*

© 2018 Palo Alto Networks, Inc.

STEP 7 | Save the portal configuration.

1. Click **OK** to save the settings and close the GlobalProtect Portal Configuration dialog.
2. **Commit** your changes.

© 2018 Palo Alto Networks, Inc.

Prepare the Satellite to Join the LSVPN

To participate in the LSVPN, the satellites require a minimal amount of configuration. Because the required configuration is minimal, you can pre-configure the satellites before shipping them to your branch offices for installation.

STEP 1 | Configure a Layer 3 Interface (see Configure Layer 3 Interfaces).

This is the physical interface the satellite will use to connect to the portal and the gateway. This interface must be in a zone that allows access outside of the local trust network. As a best practice, create a dedicated zone for VPN connections for visibility and control over traffic destined for the corporate gateways.

STEP 2 | Configure the logical tunnel interface for the tunnel to use to establish VPN tunnels with the GlobalProtect gateways.



IP addresses are not required on the tunnel interface unless you plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.

1. Select **Network** > **Interfaces** > **Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down and select an existing zone or create a separate zone for VPN tunnel traffic by clicking **New Zone** and defining a **Name** for new zone (for example *lsvpnsat*).
4. In the **Virtual Router** drop-down, select **default**.
5. (Optional) To assign an IP address to the tunnel interface:
   - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
   - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
6. To save the interface configuration, click **OK**.

STEP 3 | If you generated the portal server certificate using a Root CA that is not trusted by the satellites (for example, if you used self-signed certificates), import the root CA certificate used to issue the portal server certificate.

The root CA certificate is required to enable the satellite to establish the initial connection with the portal to obtain the LSVPN configuration.

1. Download the CA certificate that was used to generate the portal server certificates. If you are using self-signed certificates, export the root CA certificate from the portal as follows:
   1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates**.
   2. Select the CA certificate, and click **Export**.
   3. Select **Base64 Encoded Certificate (PEM)** from the **File Format** drop-down and click **OK** to download the certificate. (You do not need to export the private key.)
2. Import the root CA certificate you just exported onto each satellite as follows.
   1. Select **Device** > **Certificate Management** > **Certificates** > **Device Certificates** and click **Import**.
   2. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.
   3. **Browse** to the **Certificate File** you downloaded from the CA.
   4. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.

© 2018 Palo Alto Networks, Inc.

5. Select the certificate you just imported on the **Device Certificates** tab to open it.
6. Select **Trusted Root CA** and then click **OK**.

STEP 4 | Configure the IPSec tunnel configuration.

1. Select **Network** > **IPSec Tunnels** and click **Add**.
2. On the **General** tab, enter a descriptive **Name** for the IPSec configuration.
3. Select the **Tunnel Interface** you created for the satellite.
4. Select **GlobalProtect Satellite** as the **Type**.
5. Enter the IP address or FQDN of the portal as the **Portal Address**.
6. Select the Layer 3 **Interface** you configured for the satellite.
7. Select the **IP Address** to use on the selected interface. You can select an **IPv4** address, an **IPv6** address, or both. Specify if you want **IPv6 preferred for portal registration**.

STEP 5 | (Optional) Configure the satellite to publish local routes to the gateway.

Pushing routes to the gateway enables traffic to the subnets local to the satellite via the gateway. However, you must also configure the gateway to accept the routes as detailed in Configure GlobalProtect Gateways for LSVPN.

1. To enable the satellite to push routes to the gateway, on the **Advanced** tab select **Publish all static and connected routes to Gateway**.

   If you select this check box, the firewall will forward all static and connected routes from the satellite to the gateway. However, to prevent the creation of routing loops, the firewall will apply some route filters, such as the following:

   - Default routes
   - Routes within a virtual router other than the virtual router associated with the tunnel interface
   - Routes using the tunnel interface
   - Routes using the physical interface associated with the tunnel interface
2. (Optional) If you only want to push routes for specific subnets rather than all routes, click **Add** in the Subnet section and specify which subnet routes to publish.

STEP 6 | Save the satellite configuration.

1. Click **OK** to save the IPSec tunnel settings.
2. Click **Commit**.

STEP 7 | If required, provide the credentials to allow the satellite to authenticate to the portal.

This step is only required if the portal was unable to find a serial number match in its configuration or if the serial number didn't work. In this case, the satellite will not be able to establish the tunnel with the gateway(s).

1. Select **Network** > **IPSec Tunnels** and click the **Gateway Info** link in the Status column of the tunnel configuration you created for the LSVPN.
2. Click the **enter credentials** link in the **Portal Status** field and username and password required to authenticate the satellite to the portal.

   After the portal successfully authenticates to the portal, it will receive its signed certificate and configuration, which it will use to connect to the gateway(s). You should see the tunnel establish and the **Status** change to **Active**.

© 2018 Palo Alto Networks, Inc.

Verify the LSVPN Configuration

After configuring the portal, gateways, and satellites, verify that the satellites are able to connect to the portal and gateway and establish VPN tunnels with the gateway(s).

STEP 1 | Verify satellite connectivity with portal.

From the firewall hosting the portal, verify that satellites are successfully connecting by selecting **Network** > **GlobalProtect** > **Portal** and clicking **Satellite Info** in the Info column of the portal configuration entry.

STEP 2 | Verify satellite connectivity with the gateway(s).

On each firewall hosting a gateway, verify that satellites are able to establish VPN tunnels by selecting **Network** > **GlobalProtect** > **Gateways** and click **Satellite Info** in the Info column of the gateway configuration entry. Satellites that have successfully established tunnels with the gateway will display on the **Active Satellites** tab.

STEP 3 | Verify LSVPN tunnel status on the satellite.

On each firewall hosting a satellite, verify the tunnel status by selecting **Network** > **IPSec Tunnels** and verify active Status as indicated by a green icon.

© 2018 Palo Alto Networks, Inc.

LSVPN Quick Configs

The following sections provide step-by-step instructions for configuring some common GlobalProtect LSVPN deployments:

- Basic LSVPN Configuration with Static Routing on page 788
- Advanced LSVPN Configuration with Dynamic Routing on page 790
- Advanced LSVPN Configuration with iBGP on page 792

Basic LSVPN Configuration with Static Routing

This quick config shows the fastest way to get up and running with LSVPN. In this example, a single firewall at the corporate headquarters site is configured as both a portal and a gateway. Satellites can be quickly and easily deployed with minimal configuration for optimized scalability.



The following workflow shows the steps for setting up this basic configuration:

STEP 1 | Configure a Layer 3 interface.

In this example, the Layer 3 interface on the portal/gateway requires the following configuration:

- **Interface**—ethernet1/11
- **Security Zone**—lsvpn-tun
- **IPv4**—203.0.113.11/24

STEP 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.

To enable visibility into users and groups connecting over the VPN, enable User-ID in the zone where the VPN tunnels terminate.

In this example, the Tunnel interface on the portal/gateway requires the following configuration:

- **Interface**—tunnel.1
- **Security Zone**—lsvpn-tun

STEP 3 | Create the Security policy rule to enable traffic flow between the VPN zone where the tunnel terminates (lsvpn-tun) and the trust zone where the corporate applications reside (L3-Trust).

See Create a Security Policy Rule.

© 2018 Palo Alto Networks, Inc.

STEP 4 | Assign an SSL/TLS Service profile to the portal/gateway. The profile must reference a self-signed server certificate.

The certificate subject name must match the FQDN or IP address of the Layer 3 interface you create for the portal/gateway.

1. On the firewall hosting the GlobalProtect portal, create the root CA certificate for signing the certificates of the GlobalProtect components. In this example, the root CA certificate, **lsvpn-CA**, will be used to issue the server certificate for the portal/gateway. In addition, the portal will use this root CA certificate to sign the CSRs from the satellites.
2. Create SSL/TLS service profiles for the GlobalProtect portal and gateways.

   Because the portal and gateway are on the same interface in this example, they can share an SSL/TLS Service profile that uses the same server certificate. In this example, the profile is named **lsvpnserver**.

STEP 5 | Create a certificate profile.

In this example, the certificate profile **lsvpn-profile** references the root CA certificate **lsvpn-CA**. The gateway will use this certificate profile to authenticate satellites attempting to establish VPN tunnels.

STEP 6 | Configure an authentication profile for the portal to use if the satellite serial number is not available.

1. Create one type of server profile on the portal:
   - Add a RADIUS server profile.

     You can use RADIUS to integrate with a Multi-Factor Authentication service.

   - Add a TACACS+ server profile.
   - Add a SAML IdP server profile.
   - Add a Kerberos server profile.
   - Add an LDAP server profile. If you use LDAP to connect to Active Directory (AD), create a separate LDAP server profile for every AD domain.
2. Configure an authentication profile. In this example, the profile **lsvpn-sat** is used to authenticate satellites.

STEP 7 | Configure GlobalProtect Gateways for LSVPN.

Select **Network** > **GlobalProtect** > **Gateways** and **Add** a configuration. This example requires the following gateway configuration:

- **Interface**—ethernet1/11
- **IP Address**—203.0.113.11/24
- **SSL/TLS Server Profile**—lsvpnserver
- **Certificate Profile**—lsvpn-profile
- **Tunnel Interface**—tunnel.1
- **Primary DNS/Secondary DNS**—4.2.2.1/4.2.2.2
- **IP Pool**—2.2.2.111-2.2.2.120
- **Access Route**—10.2.10.0/24

STEP 8 | Configure the Portal.

© 2018 Palo Alto Networks, Inc.

Select **Network** > **GlobalProtect** > **Portal** and **Add** a configuration. This example requires the following portal configuration:

- **Interface**—ethernet1/11
- **IP Address**—203.0.113.11/24
- **SSL/TLS Server Profile**—lsvpnserver
- **Authentication Profile**—lsvpn-sat

STEP 9 | Define the Satellite Configurations.

On the **Satellite** tab in the portal configuration, **Add** a Satellite configuration and a Trusted Root CA and specify the CA the portal will use to issue certificates for the satellites. In this example the required settings are as following:

- **Gateway**—203.0.113.11
- **Issuing Certificate**—lsvpn-CA
- **Trusted Root CA**—lsvpn-CA

STEP 10 | Prepare the Satellite to Join the LSVPN.

The satellite configuration in this example requires the following settings:

Interface Configuration

- Layer 3 interface—ethernet1/1, 203.0.113.13/24
- Tunnel interface—tunnel.2
- Zone—lsvpnsat

Root CA Certificate from Portal

- lsvpn-CA

IPSec Tunnel Configuration

- **Tunnel Interface**—tunnel.2
- **Portal Address**—203.0.113.11
- **Interface**—ethernet1/1
- **Local IP Address**—203.0.113.13/24
- **Publish all static and connected routes to Gateway**—enabled

Advanced LSVPN Configuration with Dynamic Routing

In larger LSVPN deployments with multiple gateways and many satellites, investing a little more time in the initial configuration to set up dynamic routing will simplify the maintenance of gateway configurations because access routes will update dynamically. The following example configuration shows how to extend the basic LSVPN configuration to configure OSPF as the dynamic routing protocol.

Setting up an LSVPN to use OSPF for dynamic routing requires the following additional steps on the gateways and the satellites:

- Manual assignment of IP addresses to tunnel interfaces on all gateways and satellites.
- Configuration of OSPF point-to-multipoint (P2MP) on the virtual router on all gateways and satellites. In addition, as part of the OSPF configuration on each gateway, you must manually define the tunnel IP address of each satellite as an OSPF neighbor. Similarly, on each satellite, you must manually define the tunnel IP address of each gateway as an OSPF neighbor.

© 2018 Palo Alto Networks, Inc.

Although dynamic routing requires additional setup during the initial configuration of the LSVPN, it reduces the maintenance tasks associated with keeping routes up to date as topology changes occur on your network.

The following figure shows an LSVPN dynamic routing configuration. This example shows how to configure OSPF as the dynamic routing protocol for the VPN.



For a basic setup of a LSVPN, follow the steps in Basic LSVPN Configuration with Static Routing. You can then complete the steps in the following workflow to extend the configuration to use dynamic routing rather than static routing.

STEP 1 | Add an IP address to the tunnel interface configuration on each gateway and each satellite.

Complete the following steps on each gateway and each satellite:

1. Select **Network** > **Interfaces** > **Tunnel** and select the tunnel configuration you created for the LSVPN to open the Tunnel Interface dialog.

   If you have not yet created the tunnel interface, see Step 2 in Create Interfaces and Zones for the LSVPN.
2. On the **IPv4** tab, click **Add** and then enter an IP address and subnet mask. For example, to add an IP address for the gateway tunnel interface you would enter 2.2.2.100/24.
3. Click **OK** to save the configuration.

STEP 2 | Configure the dynamic routing protocol on the gateway.

To configure OSPF on the gateway:

1. Select **Network** > **Virtual Routers** and select the virtual router associated with your VPN interfaces.
2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.
3. If you are creating a new area, enter an **Area ID** on the **Type** tab.
4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.
5. Select **p2mp** as the **Link Type**.
6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each satellite, for example 2.2.2.111.
7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.
8. Repeat this step each time you add a new satellite to the LSVPN.

STEP 3 | Configure the dynamic routing protocol on the satellite.

To configure OSPF on the satellite:

1. Select **Network** > **Virtual Routers** and select the virtual router associated with your VPN interfaces.

© 2018 Palo Alto Networks, Inc.

2. On the **Areas** tab, click **Add** to create the backbone area, or, if it is already configured, click on the area ID to edit it.
3. If you are creating a new area, enter an **Area ID** on the **Type** tab.
4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN.
5. Select **p2mp** as the **Link Type**.
6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each GlobalProtect gateway, for example 2.2.2.100.
7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway.
8. Repeat this step each time you add a new gateway.

STEP 4 | Verify that the gateways and satellites are able to form router adjacencies.

- On each satellite and each gateway, confirm that peer adjacencies have formed and that routing table entries have been created for the peers (that is, the satellites have routes to the gateways and the gateways have routes to the satellites). Select **Network** > **Virtual Router** and click the **More Runtime Stats** link for the virtual router you are using for the LSVPN. On the Routing tab, verify that the LSVPN peer has a route.
- On the **OSPF** > **Interface** tab, verify that the **Type** is **p2mp**.
- On the **OSPF** > **Neighbor** tab, verify that the firewalls hosting your gateways have established router adjacencies with the firewalls hosting your satellites and vice versa. Also verify that the **Status** is **Full**, indicating that full adjacencies have been established.

Advanced LSVPN Configuration with iBGP

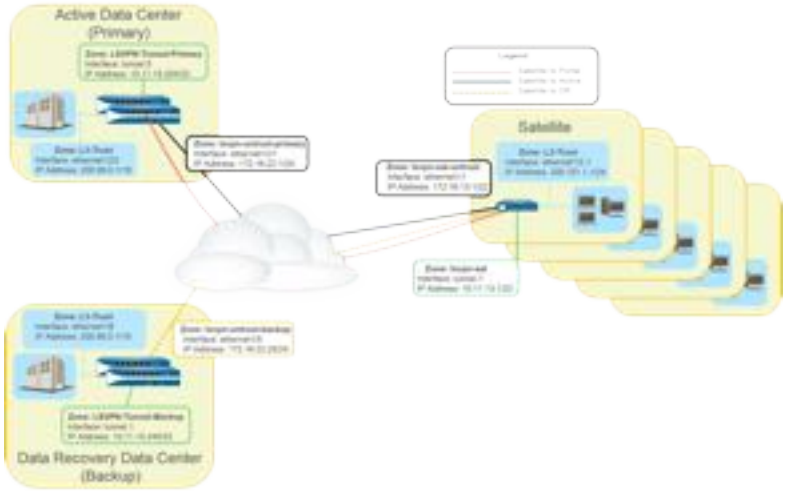
This use case illustrates how GlobalProtect LSVPN securely connects distributed office locations with primary and disaster recovery data centers that house critical applications for users and how internal border gateway protocol (iBGP) eases deployment and upkeep. Using this method, you can extend up to 500 satellite offices connecting to a single gateway.

BGP is a highly scalable, dynamic routing protocol that is ideal for hub-and-spoke deployments such as LSVPN. As a dynamic routing protocol, it eliminates much of the overhead associated with access routes (static routes) by making it relatively easy to deploy additional satellite firewalls. Due to its route filtering capabilities and features such as multiple tunable timers, route dampening, and route refresh, BGP scales to a much higher number of routing prefixes with greater stability than other routing protocols like RIP and OSPF. In the case of iBGP, a peer group, which includes all the satellites and gateways in the LSVPN deployment, establishes adjacencies over the tunnel endpoints. The protocol then implicitly takes control of route advertisements, updates, and convergence.

In this example configuration, an active/passive HA pair of PA-5050 firewalls is deployed in the primary (active) data center and acts as the portal and primary gateway. The disaster recovery data center also has two PA-5050s in an active/passive HA pair acting as the backup LSVPN gateway. The portal and gateways serve 500 PA-200s deployed as LSVPN satellites in branch offices.

Both data center sites advertise routes but with different metrics. As a result, the satellites prefer and install the active data center's routes. However, the backup routes also exist in the local routing information base (RIB). If the active data center fails, the routes advertised by that data center are removed and replaced with routes from the disaster recovery data center's routes. The failover time depends on selection of iBGP times and routing convergence associated with iBGP.

© 2018 Palo Alto Networks, Inc.



The following workflow shows the steps for configuring this deployment:

STEP 1 | Create Interfaces and Zones for the LSVPN.

Portal and Primary gateway:

- **Zone**: LSVPN-Untrust-Primary
- Interface: ethernet1/21
- IPv4: 172.16.22.1/24
- **Zone**: L3-Trust
- **Interface**: ethernet1/23
- **IPv4**: 200.99.0.1/16

Backup gateway:

- **Zone**: LSVPN-Untrust-Primary
- **Interface**: ethernet1/5
- **IPv4**: 172.16.22.25/24
- **Zone**: L3-Trust
- **Interface**: ethernet1/6
- **IPv4**: 200.99.0.1/16

Satellite:

- **Zone**: LSVPN-Sat-Untrust
- **Interface**: ethernet1/1
- **IPv4**: 172.16.13.1/22
- **Zone**: L3-Trust
- **Interface**: ethernet1/2.1
- **IPv4**: 200.101.1.1/24

Configure the zones, interfaces, and IP addresses on each satellite. The interface and local IP address will be different for each satellite. This interface is used for the VPN connection to the portal and gateway.

© 2018 Palo Alto Networks, Inc.

STEP 2 | On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.

Primary gateway:

- **Interface**: tunnel.5
- **IPv4**: 10.11.15.254/22
- **Zone**: LSVPN-Tunnel-Primary

Backup gateway:

- **Interface**: tunnel.1
- **IPv4**: 10.11.15.245/22
- **Zone**: LSVPN-Tunnel-Backup

STEP 3 | Enable SSL Between GlobalProtect LSVPN Components.

The gateway uses the self-signed root certificate authority (CA) to issue certificates for the satellites in a GlobalProtect LSVPN. Because one firewall houses the portal and primary gateway, a single certificate is used for authenticating to the satellites. The same CA is used to generate a certificate for the backup gateway. The CA generates certificates that pushed to the satellites from the portal and then used by the satellites to authenticate to the gateways.

You must also generate a certificate from the same CA for the backup gateway, allowing it to authenticate with the satellites.

1. On the firewall hosting the GlobalProtect portal, create the root CA certificate for signing the certificates of the GlobalProtect components. In this example, the root CA certificate is called CA-cert.
2. Create SSL/TLS service profiles for the GlobalProtect portal and gateways. Because the GlobalProtect portal and primary gateway are the same firewall interface, you can use the same server certificate for both components.
   - **Root CA Certificate**: CA-Cert
   - **Certificate Name**: LSVPN-Scale
3. Deploy the self-signed server certificates to the gateways.
4. Import the root CA certificate used to issue server certificates for the LSVPN components.
5. Create a certificate profile.
6. Repeat steps 2 through 5 on the backup gateway with the following settings:
   - **Root CA Certificate**: CA-cert
   - **Certificate Name**: LSVPN-back-GW-cert

STEP 4 | Configure GlobalProtect Gateways for LSVPN.

1. Select **Network** > **GlobalProtect** > **Gateways** and click **Add**.
2. On the **General** tab, name the primary gateway **`LSVPN-Scale`**.
3. Under **Network Settings**, select **`ethernet1/21`** as the primary gateway interface and enter **`172.16.22.1/24`** as the IP address.
4. On the **Authentication** tab, select the LSVPN-Scale certificate created in 3.
5. Select **Satellite** > **Tunnel Settings** and select **Tunnel Configuration**. Set the **Tunnel Interface** to tunnel.5. All satellites in this use case connect to a single gateway, so a single satellite configuration is needed. Satellites are matched based on their serial numbers, so no satellites will need to authenticate as a user.
6. On **Satellite** > **Network Settings**, define the pool of IP address to assign to the tunnel interface on the satellite once the VPN connection is established. Because this use case uses dynamic routing, the Access Routes setting remains blank.
7. Repeat steps 1 through 5 on the backup gateway with the following settings:

© 2018 Palo Alto Networks, Inc.

- **Name**: LSVPN-backup
- **Gateway interface**: ethernet1/5
- **Gateway IP**: 172.16.22.25/24
- **Server cert**: LSVPN-backup-GW-cert
- **Tunnel interface**: tunnel.1

STEP 5 | Configure iBGP on the primary and backup gateways and add a redistribution profile to allow the satellites to inject local routes back to the gateways.

Each satellite office manages its own network and firewall, so the redistribution profile called ToAllSat is configured to redistribute local routes back to the GlobalProtect gateway.

1. Select **Network** > **Virtual Routers** and **Add** a virtual router.
2. On **Router Settings**, add the **Name** and **Interface** for the virtual router.
3. On **Redistribution Profile** and select **Add**.
    1. Name the redistribution profile **`ToAllSat`** and set the **Priority** to 1.
    2. Set Redistribute to **Redist**.
    3. **Add `ethernet1/23`** from the Interface drop-down.
    4. Click **OK**.
4. Select **BGP** on the Virtual Router to configure BGP.
    1. On **BGP** > **General**, select **Enable**.
    2. Enter the gateway IP address as the **Router ID** (**`172.16.22.1`**) and **`1000`** as the **AS Number**.
    3. In the Options section, select **Install Route**.
    4. On **BGP** > **Peer Group**, click **Add** a peer group with all the satellites that will connect to the gateway.
    5. On **BGP** > **Redist Rules**, **Add** the **`ToAllSat`** redistribution profile you created previously.
5. Click **OK**.
6. Repeat steps 1 through 5 on the backup gateway using **`ethernet1/6`** for the redistribution profile.

STEP 6 | Prepare the Satellite to Join the LSVPN.

The configuration shown is a sample of a single satellite.

Repeat this configuration each time you add a new satellite to the LSVPN deployment.

1. Configure a tunnel interface as the tunnel endpoint for the VPN connection to the gateways.
2. Set the IPSec tunnel type to GlobalProtect Satellite and enter the IP address of the GlobalProtect Portal.
3. Select **Network** > **Virtual Routers** and **Add** a virtual router.
4. On **Router Settings**, add the **Name** and **Interface** for the virtual router.
5. Select **Virtual Router** > **Redistribution Profile** and **Add** a profile with the following settings.
    1. Name the redistribution profile **`ToLSVPNGW`** and set the **Priority** to 1.
    2. **Add** an **Interface `ethernet1/2.1`**.
    3. Click **OK**.
6. Select **BGP** > **General**, **Enable** BGP and configure the protocol as follows:
    1. Enter the gateway IP address as the **Router ID** (**`172.16.22.1`**) and **`1000`** as the **AS Number**.
    2. In the Options section, select **Install Route**.
    3. On **BGP** > **Peer Group**, **Add** a peer group containing all the satellites that will connect to the gateway.
    4. On **BGP** > **Redist Rules**, **Add** the **`ToLSVPNGW`** redistribution profile you created previously.
7. Click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 7 | Configure the GlobalProtect Portal for LSVPN.

Both data centers advertise their routes but with different routing priorities to ensure that the active data center is the preferred gateway.

1. Select **Network** > **GlobalProtect** > **Portals** and click **Add**.
2. On **General**, enter **`LSVPN-Portal`** as the portal name.
3. On **Network Settings**, select **`ethernet1/21`** as the **Interface** and select **`172.16.22.1/24`** as the **IP Address**.
4. On the **Authentication** tab, select the previously created primary gateway SSL/TLS Profile **`LSVPN-Scale`** from the **SSL/TLS Service Profile** drop-down menu.
5. On the **Satellite** tab, **Add** a satellite and **Name** it **`sat-config-1`**.
6. Set the **Configuration Refresh Interval** to **`12`**.
7. On **GlobalProtect Satellite** > **Devices**, add the serial number and hostname of each satellite device in the LSVPN.
8. On **GlobalProtect Satellite** > **Gateways**, add the name and IP address of each gateway. Set the routing priority of the primary gateway to 1 and the backup gateway to 10 to ensure that the active data center is the preferred gateway.

STEP 8 | Verify the LSVPN Configuration.

STEP 9 | (Optional) Add a new site to the LSVPN deployment.

1. Select **Network** > **GlobalProtect** > **Portals** > **GlobalProtect Portal** > **Satellite Configuration** > **GlobalProtect Satellite** > **Devices** to add the serial number of the new satellite to the GlobalProtect portal.
2. Configure the IPSec tunnel on the satellite with the GlobalProtect Portal IP address.
3. Select **Network** > **Virtual Router** > **BGP** > **Peer Group** to add the satellite to the BGP Peer Group configuration on each gateway.
4. Select **Network** > **Virtual Router** > **BGP** > **Peer Group** to add the gateways to the BGP Peer Group configuration on the new satellite.

Networking

All Palo Alto Networks® next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, and enables you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports.

The following topics describe networking concepts and how to integrate Palo Alto Networks next-generation firewalls into your network.

> Configure Interfaces
> Virtual Routers
> Service Routes
> Static Routes
> RIP
> OSPF
> BGP
> IP Multicast on page 857
> Route Redistribution
> DHCP
> DNS
> NAT
> NPTv6
> NAT64
> ECMP
> LLDP
> BFD
> Session Settings and Timeouts
> Tunnel Content Inspection

© 2018 Palo Alto Networks, Inc.

Configure Interfaces

A Palo Alto Networks next-generation firewall can operate in multiple deployments at once because the deployments occur at the interface level. For example, you can configure some interfaces for Layer 3 interfaces to integrate the firewall into your dynamic routing environment, while configuring other interfaces to integrate into your Layer 2 switching network. The following topics describe each type of interface deployment and how to configure the corresponding interface types:

- Tap Interfaces
- Virtual Wire Interfaces
- Layer 2 Interfaces
- Layer 3 Interfaces
- Configure an Aggregate Interface Group
- Use Interface Management Profiles to Restrict Access

Tap Interfaces

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

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



When deployed in tap mode, the firewall is not able to take action, such as block traffic or apply QoS traffic control.

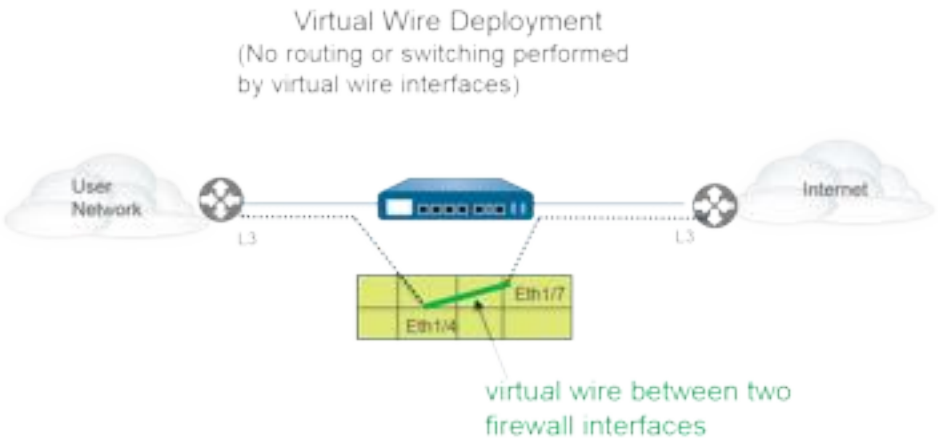
Virtual Wire Interfaces

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; hence, 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 the two connected interfaces on the firewall need not do any switching or routing. For these two interfaces, the firewall is considered a *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 surrounding network devices. The virtual wire supports blocking or allowing traffic based on virtual LAN (VLAN) tags, in addition to supporting security policy rules, App-ID, Content-ID, User-ID, decryption, LLDP, active/passive and active/active HA, QoS, zone protection (with some exceptions), non-IP protocol protection, DoS protection, packet buffer protection, tunnel content inspection, and NAT.

© 2018 Palo Alto Networks, Inc.



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 one of the virtual wire interfaces receives a frame or packet, it ignores any Layer 2 or Layer 3 addresses for switching or routing purposes, but applies your 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 wouldn't use a virtual wire deployment for interfaces that need to support switching, VPN tunnels, or routing because they require a Layer 2 or Layer 3 address. A virtual wire interface doesn't use an interface management profile, which controls services such as HTTP and ping and therefore requires the interface have an IP address.

All firewalls shipped from the factory have two Ethernet ports (ports 1 and 2) preconfigured as virtual wire interfaces, and these interfaces allow all untagged traffic.

If you don't intend to use the preconfigured virtual wire, you must delete that configuration to prevent it from interfering with other settings you configure on the firewall. See Set Up Network Access for External Services.

- Layer 2 and Layer 3 Packets over a Virtual Wire
- Port Speeds of Virtual Wire Interfaces
- LLDP over a Virtual Wire
- Aggregated Interfaces for a Virtual Wire
- Virtual Wire Support of High Availability
- Zone Protection for a Virtual Wire Interface
- VLAN-Tagged Traffic
- Virtual Wire Subinterfaces
- Configure Virtual Wires

Layer 2 and Layer 3 Packets over a Virtual Wire

A virtual wire interface will allow Layer 2 and Layer 3 packets from connected devices to pass transparently as long as the policies applied to the zone or interface allow the traffic. The virtual wire interfaces themselves don't participate in routing or switching.

For example, the firewall doesn't decrement the TTL in a traceroute packet going over the virtual link because the link is transparent and doesn't count as a hop. Packets such as Operations, Administration and Maintenance (OAM) protocol data units (PDUs), for example, don't terminate at the firewall. Thus, the virtual wire allows the firewall to maintain a transparent presence acting as a pass-through link, while still providing security, NAT, and QoS services.

© 2018 Palo Alto Networks, Inc.

In order for bridge protocol data units (BPDUs) and other Layer 2 control packets (which are typically untagged) to pass through a virtual wire, the interfaces must be attached to a virtual wire object that allows untagged traffic, and that is the default. If the virtual wire object **Tag Allowed** field is empty, the virtual wire allows untagged traffic. (Security policy rules don't apply to Layer 2 packets.)

In order for routing (Layer 3) control packets to pass through a virtual wire, you must apply a security policy rule that allows the traffic to pass through. For example, apply a security policy rule that allows an application such as BGP or OSPF.

If you want to be able to apply security policy rules to a zone for IPv6 traffic arriving at a virtual wire interface on the firewall, enable IPv6 firewalling. Otherwise, IPv6 traffic is forwarded transparently across the wire.

If you enable multicast firewalling for a virtual wire object and apply it to a virtual wire interface, the firewall inspects multicast traffic and forwards it or not, based on security policy rules. If you don't enable multicast firewalling, the firewall simply forwards multicast traffic transparently.

Fragmentation on a virtual wire occurs the same as in other interface deployment modes.

Port Speeds of Virtual Wire Interfaces

Different firewall models provide various numbers of copper and fiber optic ports, which operate at different speeds. A virtual wire can bind two Ethernet ports of the same type (both copper or both fiber optic), or bind a copper port with a fiber optic port. By default, the **Link Speed** of copper ports on the firewall is set to **auto**, which means the firewall automatically negotiates their speed and transmission mode (**Link Duplex**). When you Configure Virtual Wires, you can also select a specific **Link Speed** and **Link Duplex** but the values for these settings must be the same for both ports in any single virtual wire.

LLDP over a Virtual Wire

Virtual wire interfaces can use LLDP to discover neighboring devices and their capabilities, and LLDP allows neighboring devices to detect the presence of the firewall in the network. LLDP makes troubleshooting easier especially on a virtual wire, where the firewall would typically go undetected by a ping or traceroute passing through the virtual wire. LLDP provides a way for other devices to detect the firewall in the network. Without LLDP, it is practically impossible for network management systems to detect the presence of a firewall through the virtual link.

Aggregated Interfaces for a Virtual Wire

You can Configure an Aggregate Interface Group of virtual wire interfaces, but virtual wires don't use LACP. If you configure LACP on devices that connect the firewall to other networks, the virtual wire will pass LACP packets transparently without performing LACP functions.



In order for aggregate interface groups to function properly, ensure all links belonging to the same LACP group on the same side of the virtual wire are assigned to the same zone.

Virtual Wire Support of High Availability

If you configure the firewall to perform path monitoring for High Availability using a virtual wire path group, the firewall attempts to resolve ARP for the configured destination IP address by sending ARP packets out both of the virtual wire interfaces. The destination IP address that you are monitoring must be on the same subnetwork as one of the devices surrounding the virtual wire.

Virtual wire interfaces support both active/passive and active/active HA. For an active/active HA deployment with a virtual wire, the scanned packets must be returned to the receiving firewall to preserve the forwarding path. Therefore, if a firewall receives a packet that belongs to the session that the peer HA firewall owns, it sends the packet across the HA3 link to the peer.

© 2018 Palo Alto Networks, Inc.

For PAN-OS 7.1 and later releases, you can configure the passive firewall in an HA pair to allow peer devices on either side of the firewall to pre-negotiate LLDP and LACP over a virtual wire before an HA failover occurs. Such a configuration for LACP and LLDP Pre-Negotiation for Active/Passive HA speeds up HA failovers.

Zone Protection for a Virtual Wire Interface

You can apply zone protection to a virtual wire interface, but because virtual wire interfaces don't perform routing, you can't apply Packet-Based Attack Protection to packets coming with a spoofed IP address, nor can you suppress ICMP TTL Expired error packets or ICMP Frag Needed packets.

By default, a virtual wire interface forwards all non-IP traffic it receives. However, you can apply a zone protection profile with Protocol Protection to block or allow certain non-IP protocol packets between security zones on a virtual wire.

VLAN-Tagged Traffic

Virtual wire interfaces by default allow all untagged traffic. You can, however, use a virtual wire to connect two interfaces and configure either interface to block or allow traffic based on the virtual LAN (VLAN) tags. VLAN tag 0 indicates untagged traffic.

You can also create multiple subinterfaces, add them into different zones, and then classify traffic according to a VLAN tag or a combination of a VLAN tag with IP classifiers (address, range, or subnet) to apply granular policy control for specific VLAN tags or for VLAN tags from a specific source IP address, range, or subnet.

Virtual Wire Subinterfaces

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

- **VLAN tags**—The example in Virtual Wire Deployment with Subinterfaces (VLAN Tags only) shows an ISP using virtual wire subinterfaces with VLAN tags to separate traffic for two different customers.
- **VLAN tags in conjunction with IP classifiers (address, range, or subnet)**—The 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 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, and assign these interfaces to a virtual wire.
• Create subinterfaces on the parent Virtual Wire to separate CustomerA and CustomerB traffic. Make sure that the VLAN tags defined on each pair of subinterfaces that are configured as virtual wire(s) are identical. This is essential because a virtual wire does not switch VLAN tags.
• Create new subinterfaces and define IP classifiers. This task is optional and only required if you wish to add additional subinterfaces with IP classifiers for further managing traffic from a customer based on the combination of VLAN tags and a specific source IP address, range or subnet. You can also use IP classifiers for managing untagged traffic. To do so, you must create a sub-interface with the vlan tag "0", and define subinterface(s) with IP classifiers for managing untagged traffic using IP classifiers.

© 2018 Palo Alto Networks, Inc.

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



Figure 11: Virtual Wire Deployment with Subinterfaces (VLAN Tags only)

Virtual Wire Deployment with Subinterfaces (VLAN Tags only) depicts CustomerA and CustomerB connected to the firewall through one physical interface, ethernet1/1, configured as a Virtual Wire; it is the ingress interface. A second physical interface, ethernet1/2, is also part of the Virtual Wire; it is the egress interface that provides access to the internet.

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

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

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

Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers) depicts CustomerA and CustomerB connected to one physical firewall that has two virtual systems (vsys), in addition to the default virtual system (vsys1). Each virtual system is an independent virtual firewall that is managed separately for each customer. Each vsys has attached interfaces/subinterfaces and security zones that are managed independently.

© 2018 Palo Alto Networks, Inc.

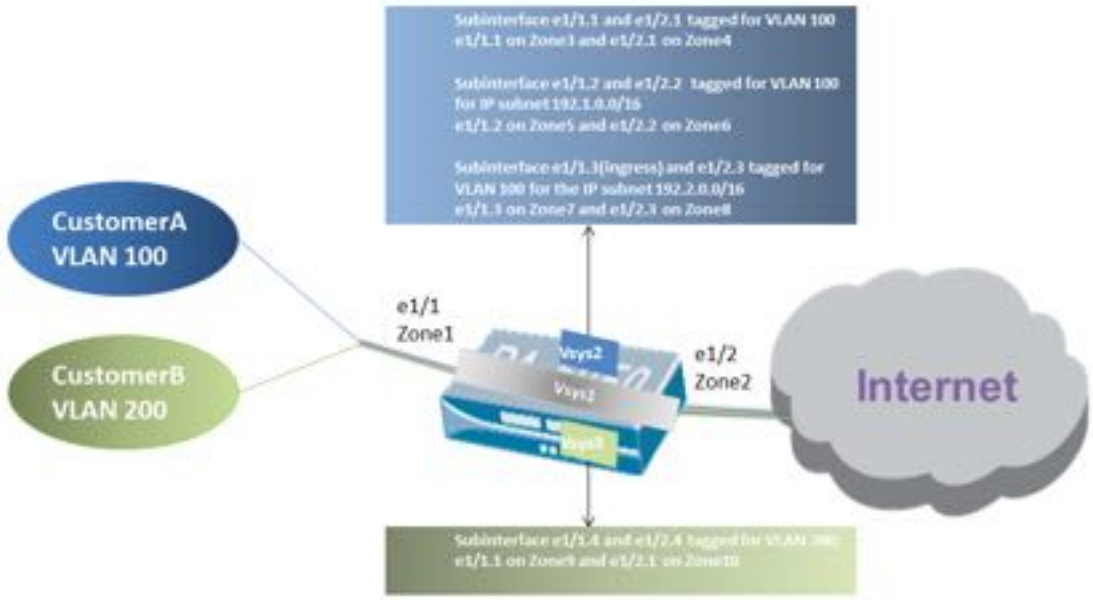


Figure 12: Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers)

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

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

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

© 2018 Palo Alto Networks, Inc.

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

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

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

Configure Virtual Wires

The following task shows how to configure two Virtual Wire Interfaces (Ethernet 1/3 and Ethernet 1/4 in this example) to create a virtual wire. The two interfaces must have the same **Link Speed** and transmission mode (**Link Duplex**). For example, a full-duplex 1000Mbps copper port matches a full-duplex 1Gbps fiber optic port.

STEP 1 | Create the first virtual wire interface.

1. Select **Network** > **Interfaces** > **Ethernet** and select an interface you have cabled (**ethernet1/3** in this example).
2. Set the **Interface Type** to **Virtual Wire**.

STEP 2 | Attach the interface to a virtual wire object.

1. While still on the same Ethernet interface, on the **Config** tab, select **Virtual Wire** and click **New Virtual Wire**.
2. Enter a **Name** for the virtual wire.
3. For **Interface1**, select the interface you just configured (**ethernet1/3**). (Only interfaces configured as virtual wire interfaces appear in the drop-down.)
4. For **Tag Allowed**, enter **0** to indicate untagged traffic (such as BPDUs and other Layer 2 control traffic) is allowed. The absence of a tag implies tag 0. Enter additional allowed tag integers or ranges of tags, separated by commas (default is 0; range is 0 to 4,094).
5. Select **Multicast Firewalling** if you want to be able to apply security policy rules to multicast traffic going across the virtual wire. Otherwise, multicast traffic is transparently forwarded across the virtual wire.
6. Select **Link State Pass Through** so the firewall can function transparently. When the firewall detects a link down state for a link of the virtual wire, it brings down the other interface in the virtual wire pair. Thus, devices on both sides of the firewall see a consistent link state, as if there were no firewall between them. If you don't select this option, link status is not propagated across the virtual wire.
7. Click **OK** to save the virtual wire object.

STEP 3 | Determine the link speed of the virtual wire interface.

1. While still on the same Ethernet interface, select **Advanced** and note or change the **Link Speed**. The port type determines the speed settings available in the drop-down. By default, copper ports are set to **auto** negotiate link speed. Both virtual wire interfaces must have the same link speed.
2. Click **OK** to save the Ethernet interface.

STEP 4 | Configure the second virtual wire interface (**ethernet1/4** in this example) by repeating the preceding steps.

© 2018 Palo Alto Networks, Inc.

When you select the **Virtual Wire** object you created, the firewall automatically adds the second virtual wire interface as **Interface2**.

STEP 5 | Create a separate security zone for each virtual wire interface.

1. Select **Network** > **Zones** and **Add** a zone.
2. Enter the **Name** of the zone (such as **`internet`**).
3. For **Location**, select the virtual system where the zone applies.
4. For **Type**, select **Virtual Wire**.
5. **Add** the **Interface** that belongs to the zone.
6. Click **OK**.

STEP 6 | (Optional) Create security policy rules to allow Layer 3 traffic to pass through.

To allow Layer 3 traffic across the virtual wire, Create a Security Policy Rule to allow traffic from the user zone to the internet zone, and another to allow traffic from the internet zone to the user zone, selecting the applications you want to allow, such as BGP or OSPF.

STEP 7 | (Optional) Enable IPv6 firewalling.

If you want to be able to apply security policy rules to IPv6 traffic arriving at the virtual wire interface, enable IPv6 firewalling. Otherwise, IPv6 traffic is forwarded transparently.

1. Select **Device** > **Setup** > **Session** and edit Session Settings.
2. Select **Enable IPv6 Firewalling**.
3. Click **OK**.

STEP 8 | **Commit** your changes.

STEP 9 | (Optional) Configure an LLDP profile and apply it to the virtual wire interfaces (see Configure LLDP).

STEP 10 | (Optional) Apply non-IP protocol control to the virtual wire zones (see Configure Protocol Protection). Otherwise, all non-IP traffic is forwarded over the virtual wire.

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 MAC address identified in the frame. Configure a Layer 2 Interface when switching is required.



Figure 13: Layer 2 Deployment

In a Layer 2 deployment, the firewall rewrites the inbound Port VLAN ID (PVID) number in a Cisco per-VLAN spanning tree (PVST+) or Rapid PVST+ bridge protocol data unit (BPDU) to the proper outbound VLAN ID number and forwards it out. The firewall rewrites such BPDUs on Layer 2 Ethernet and Aggregated Ethernet (AE) interfaces only.

A Cisco switch must have the loopguard disabled for the PVST+ or Rapid PVST+ BPDU rewrite to function properly on the firewall.

© 2018 Palo Alto Networks, Inc.

The following topics describe the different types of Layer 2 interfaces you can configure for each type of deployment you need, including details on using virtual LANs (VLANs) for traffic and policy separation among groups.

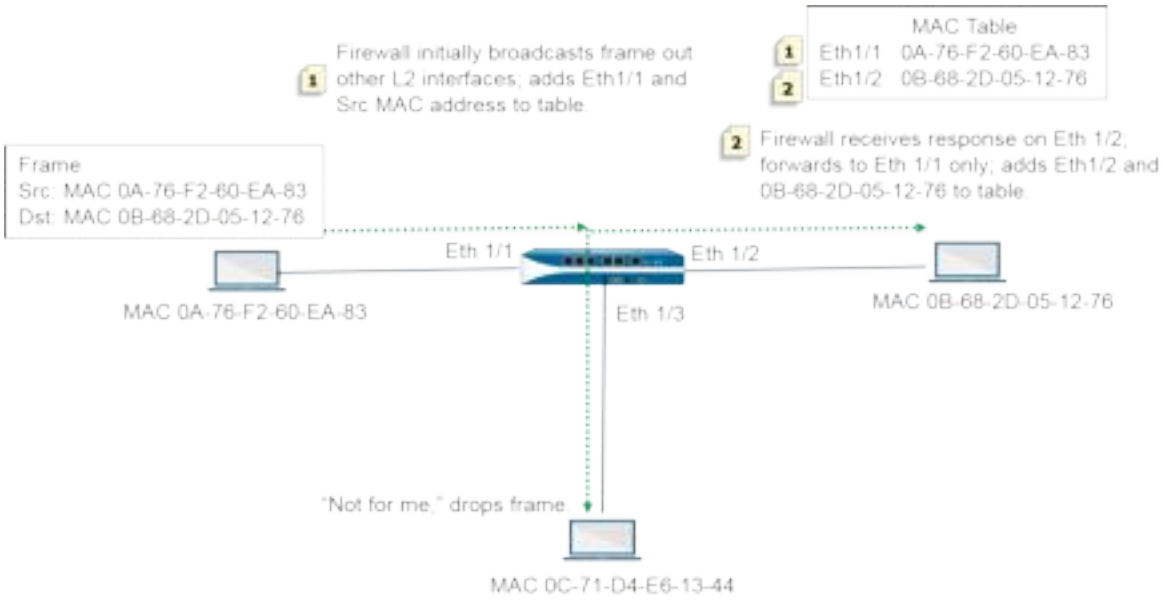
- Layer 2 Interfaces with No VLANs
- Layer 2 Interfaces with VLANs
- Configure a Layer 2 Interface
- Configure a Layer 2 Interface, Subinterface, and VLAN

Layer 2 Interfaces with No VLANs

Configure a Layer 2 Interface on the firewall so it can act as a switch in your layer 2 network (not at the edge of the network). The Layer 2 hosts are probably geographically close to each other and belong to a single broadcast domain. The firewall provides security between the Layer 2 hosts when you assign the interfaces to security zones and apply security rules to the zones.

The hosts communicate with the firewall and each other at Layer 2 of the OSI model by exchanging frames. A frame contains an Ethernet header that includes a source and destination Media Access Control (MAC) address, which is a physical hardware address. MAC addresses are 48-bit hexadecimal numbers formatted as six octets separated by a colon or hyphen (for example, 00-85-7E-46-F1-B2).

The following figure has a firewall with three Layer 2 interfaces that each connect to a Layer 2 host in a one-to-one mapping.



The firewall begins with an empty MAC table. When the host with source address 0A-76-F2-60-EA-83 sends a frame to the firewall, the firewall doesn't have destination address 0B-68-2D-05-12-76 in its MAC table, so it doesn't know which interface to forward the frame to; it broadcasts the frame to all of its Layer 2 interfaces. The firewall puts source address 0A-76-F2-60-EA-83 and associated Eth1/1 into its MAC table.

The host at 0C-71-D4-E6-13-44 receives the broadcast, but the destination MAC address is not its own MAC address, so it drops the frame.

The receiving interface Ethernet 1/2 forwards the frame to its host. When host 0B-68-2D-05-12-76 responds, it uses the destination address 0A-76-F2-60-EA-83, and the firewall adds to its MAC table Ethernet 1/2 as the interface to reach 0B-68-2D-05-12-76.

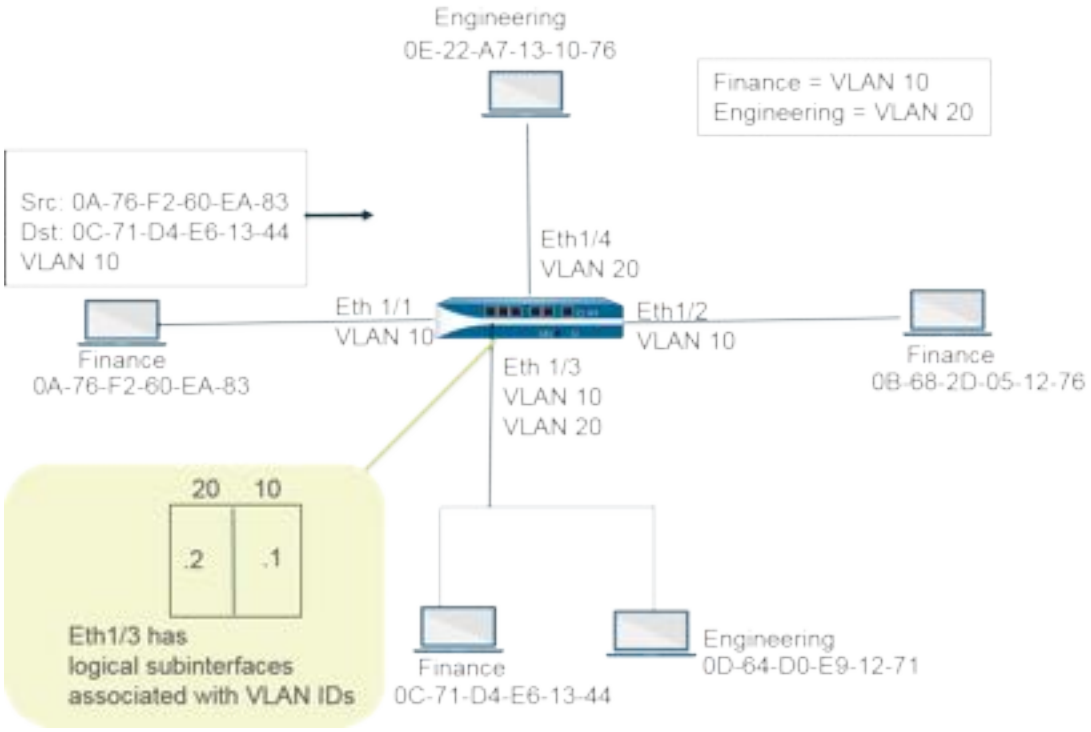
© 2018 Palo Alto Networks, Inc.

Layer 2 Interfaces with VLANs

When your organization wants to divide a LAN into separate virtual LANs (VLANs) to keep traffic and policies for different departments separate, you can logically group Layer 2 hosts into VLANs and thus divide a Layer 2 network segment into broadcast domains. For example, you can create VLANs for the Finance and Engineering departments. To do so, Configure a Layer 2 Interface, Subinterface, and VLAN.

The firewall acts as a switch to forward a frame with an Ethernet header containing a VLAN ID, and the destination interface must have a subinterface with that VLAN ID in order to receive that frame and forward it to the host. You configure a Layer 2 interface on the firewall and configure one or more logical subinterfaces for the interface, each with a VLAN tag (ID).

In the following figure, the firewall has four Layer 2 interfaces that connect to Layer 2 hosts belonging to different departments within an organization. Ethernet interface 1/3 is configured with subinterface .1 (tagged with VLAN 10) and subinterface .2 (tagged with VLAN 20), thus there are two broadcast domains on that segment. Hosts in VLAN 10 belong to Finance; hosts in VLAN 20 belong to Engineering.



In this example, the host at MAC address 0A-76-F2-60-EA-83 sends a frame with VLAN ID 10 to the firewall, which the firewall broadcasts to its other L2 interfaces. Ethernet interface 1/3 accepts the frame because it's connected to the host with destination 0C-71-D4-E6-13-44 and its subinterface .1 is assigned VLAN 10. Ethernet interface 1/3 forwards the frame to the Finance host.

Configure a Layer 2 Interface

Configure Layer 2 Interfaces with No VLANs when you want Layer 2 switching and you don't need to separate traffic among VLANs.

STEP 1 | Configure a Layer 2 interface.

1. Select **Network** > **Interfaces** > **Ethernet** and select an interface. The **Interface Name** is fixed, such as ethernet1/1.
2. For **Interface Type**, select **Layer2**.
3. Select the **Config** tab and assign the interface to a **Security Zone** or create a **New Zone**.

© 2018 Palo Alto Networks, Inc.

4. Configure additional Layer 2 interfaces on the firewall that connect to other Layer 2 hosts.

STEP 2 | Commit.

Click **OK** and **Commit**.

Configure a Layer 2 Interface, Subinterface, and VLAN

Configure Layer 2 Interfaces with VLANs when you want Layer 2 switching and traffic separation among VLANs. You can optionally control non-IP protocols between security zones on a Layer 2 interface or between interfaces within a single zone on a Layer 2 VLAN.

STEP 1 | Configure a Layer 2 interface and subinterface and assign a VLAN ID.

1. Select **Network** > **Interfaces** > **Ethernet** and select an interface. The **Interface Name** is fixed, such as ethernet1/1.
2. For **Interface Type**, select **Layer2**.
3. Select the **Config** tab.
4. For **VLAN**, leave the setting **None**.
5. Assign the interface to a **Security Zone** or create a **New Zone**.
6. Click **OK**.
7. With the Ethernet interface highlighted, click **Add Subinterface**.
8. The **Interface Name** remains fixed. After the period, enter the subinterface number, in the range 1-9,999.
9. Enter a VLAN **Tag** ID in the range 1-4,094.
10. Assign the subinterface to a **Security Zone**.
11. Click **OK**.

STEP 2 | Commit.

Click **Commit**.

STEP 3 | (Optional) Apply a Zone Protection profile with protocol protection to control non-IP protocol packets between Layer 2 zones (or between interfaces within a Layer 2 zone).

Configure Protocol Protection.

Layer 3 Interfaces

In a Layer 3 deployment, the firewall routes traffic between multiple ports. Before you can Configure 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.



Figure 14: Layer 3 Deployment

The following topics describe how to configure Layer 3 interfaces, and how to use Neighbor Discovery Protocol (NDP) to provision IPv6 hosts and view the IPv6 addresses of devices on the link local network to quickly locate devices.

- Configure Layer 3 Interfaces
- Manage IPv6 Hosts Using NDP

© 2018 Palo Alto Networks, Inc.

Configure Layer 3 Interfaces

The following procedure is required to configure Layer 3 Interfaces (Ethernet, VLAN, loopback, and tunnel interfaces) with IPv4 or IPv6 addresses so that the firewall can perform routing on these interfaces. If a tunnel is used for routing or if tunnel monitoring is turned on, the tunnel needs an IP address. Before performing the following task, define one or more Virtual Routers.

You would typically use the following procedure to configure an external interface that connects to the internet and an interface for your internal network. You can configure both IPv4 and IPv6 addresses on a single interface.



PAN-OS firewall models support a maximum of 16,000 IP addresses assigned to physical or virtual Layer 3 interfaces; this maximum includes both IPv4 and IPv6 addresses.

If you're using IPv6 routes, you can configure the firewall to provide IPv6 Router Advertisements for DNS Configuration. The firewall provisions IPv6 DNS clients with Recursive DNS Server (RDNS) addresses and a DNS Search List so that the client can resolve its IPv6 DNS requests. Thus the firewall is acting like a DHCPv6 server for you.

STEP 1 | Select an interface and configure it with a security zone.

1. Select **Network** > **Interfaces** and either **Ethernet**, **VLAN**, **loopback**, or **Tunnel**, depending on what type of interface you want.
2. Select the interface to configure.
3. Select the **Interface Type**—**Layer3**.
4. On the **Config** tab, for **Virtual Router**, select the virtual router you are configuring, such as **default**.
5. For **Virtual System**, select the virtual system you are configuring if on a multi-virtual system firewall.
6. For **Security Zone**, select the zone to which the interface belongs or create a **New Zone**.
7. Click **OK**.

STEP 2 | Configure an interface with an IPv4 address.

There are three ways to assign an IPv4 address to a Layer 3 interface:

- Static
- DHCP Client—The firewall interface acts as a DHCP client and receives a dynamically assigned IP address. The firewall also provides the capability to propagate settings received by the DHCP client interface into a DHCP server operating on the firewall. This is most commonly used to propagate DNS server settings from an Internet service provider to client machines operating on the network protected by the firewall.
- PPPoE—Configure the interface as a Point-to-Point Protocol over Ethernet (PPPoE) termination point to support connectivity in a Digital Subscriber Line (DSL) environment where there is a DSL modem but no other PPPoE device to terminate the connection.

1. Select **Network** > **Interfaces** and either **Ethernet**, **VLAN**, **loopback**, or **Tunnel**, depending on what type of interface you want.
2. Select the interface to configure.
3. To configure the interface with a static IPv4 address, on the **IPv4** tab, set **Type** to **Static**.
4. **Add** a **Name** and optional **Description** for the address.
5. For **Type**, select one of the following:
   - **IP Netmask**—Enter the IP address and network mask to assign to the interface, for example, 208.80.56.100/24.

© 2018 Palo Alto Networks, Inc.

If you're using a /31 subnet mask for the Layer 3 interface address, the interface must be configured with the .1/31 address in order for utilities such as ping to work properly.

- **IP Range**—Enter an IP address range, such as 192.168.2.1-192.168.2.4.
- **FQDN**—Enter a Fully Qualified Domain Name.

6. Select **Tags** to apply to the address.
7. Click **OK**.

STEP 3 | Configure an interface with Point-to-Point Protocol over Ethernet (PPPoE). See Layer 3 Interfaces.

PPPoE is not supported in HA active/active mode.

1. Select **Network** > **Interfaces** and either **Ethernet**, **VLAN**, **loopback**, or **Tunnel**.
2. Select the interface to configure.
3. On the **IPv4** tab, set **Type** to **PPPoE**.
4. On the **General** tab, select **Enable** to activate the interface for PPPoE termination.
5. Enter the **Username** for the point-to-point connection.
6. Enter the **Password** for the username and **Confirm Password**.
7. Click **OK**.

STEP 4 | Configure an interface as a DHCP Client so that it receives a dynamically-assigned IPv4 address.

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

1. Select **Network** > **Interfaces** and either **Ethernet**, **VLAN**, **loopback**, or **Tunnel**.
2. Select the interface to configure.
3. On the **IPv4** tab, set **Type** to **DHCP Client**.
4. Select **Enable** to activate the DHCP client on the interface.
5. Select **Automatically create default route pointing to default gateway provided by server** to automatically create a default route that points to the default gateway that the DHCP server provides.
6. (Optional) Enter a **Default Route Metric** (priority level) for the default route, which the firewall uses for path selection (range is 1-65,535; no default). The lower the value, the higher the priority level.
7. Click **OK**.

STEP 5 | Configure an interface with a static IPv6 address.

1. Select **Network** > **Interfaces** and either **Ethernet**, **VLAN**, **loopback**, or **Tunnel**.
2. Select the interface to configure.
3. On the **IPv6** tab, select **Enable IPv6 on the interface** to enable IPv6 addressing on the interface.
4. For **Interface ID**, enter the 64-bit extended unique identifier (EUI-64) in hexadecimal format (for example, 00:26:08:FF:FE:DE:4E:29). If you leave this field blank, the firewall uses the EUI-64 generated from the MAC address of the physical interface. If you enable the **Use interface ID as host portion** option when adding an address, the firewall uses the Interface ID as the host portion of that address.
5. **Add** the IPv6 **Address** or select an address group.
6. Select **Enable address on interface** to enable this IPv6 address on the interface.

© 2018 Palo Alto Networks, Inc.

7. Select **Use interface ID as host portion** to use the Interface ID as the host portion of the IPv6 address.
8. (Optional) Select **Anycast** to make the IPv6 address (route) an Anycast address (route), which means multiple locations can advertise the same prefix, and IPv6 sends the anycast traffic to the node it considers the nearest, based on routing protocol costs and other factors.
9. (Ethernet interface only) Select **Send Router Advertisement** (RA) to enable the firewall to send this address in Router Advertisements, in which case you must also enable the global **Enable Router Advertisement** option on the interface (next step).
10. (Ethernet interface only) Enter the **Valid Lifetime (sec)**, in seconds, that the firewall considers the address valid. The Valid Lifetime must equal or exceed the **Preferred Lifetime (sec)** (default is 2,592,000).
11. (Ethernet interface only) Enter the **Preferred Lifetime (sec)** (in seconds) that the valid address is preferred, which means the firewall can use it to send and received traffic. After the Preferred Lifetime expires, the firewall can't use the address to establish new connections, but any existing connections are valid until the **Valid Lifetime** expires (default is 604,800).
12. (Ethernet interface only) Select **On-link** if systems that have addresses within the prefix are reachable without a router.
13. (Ethernet interface only) Select **Autonomous** if systems can independently create an IP address by combining the advertised prefix with an Interface ID.
14. Click **OK**.

STEP 6 | (Ethernet or VLAN interface using IPv6 address only) Enable the firewall to send IPv6 Router Advertisements (RAs) from an interface, and optionally tune RA parameters.



Tune RA parameters for either of these reasons: To interoperate with a router/host that uses different values. To achieve fast convergence when multiple gateways are present. For example, set lower Min Interval, Max Interval, and Router Lifetime values so the IPv6 client/host can quickly change the default gateway after the primary gateway fails, and start forwarding to another default gateway in the network.

1. Select **Network** > **Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you want to configure.
3. Select **IPv6**.
4. Select **Enable IPv6 on the interface**.
5. On the **Router Advertisement** tab, select **Enable Router Advertisement** (default is disabled).
6. (Optional) Set **Min Interval (sec)**, the minimum interval, in seconds, between RAs the firewall sends (range is 3-1,350; default is 200). The firewall sends RAs at random intervals between the minimum and maximum values you set.
7. (Optional) Set **Max Interval (sec)**, the maximum interval, in seconds, between RAs the firewall sends (range is 4-1,800; default is 600). The firewall sends RAs at random intervals between the minimum and maximum values you set.
8. (Optional) Set **Hop Limit** to apply to clients for outgoing packets (range is 1-255; default is 64). Enter 0 for no hop limit.
9. (Optional) Set **Link MTU**, the link maximum transmission unit (MTU) to apply to clients (range is 1,280-9,192; default is **unspecified**). Select **unspecified** for no link MTU.
10. (Optional) Set **Reachable Time (ms)**, the reachable time, in milliseconds, that the client will use to assume a neighbor is reachable after receiving a Reachability Confirmation message. Select **unspecified** for no reachable time value (range is 0-3,600,000; default is **unspecified**).
11. (Optional) Set **Retrans Time (ms)**, the retransmission timer that determines how long the client will wait, in milliseconds, before retransmitting Neighbor Solicitation messages. Select **unspecified** for no retransmission time (range is 0-4,294,967,295; default is **unspecified**).

© 2018 Palo Alto Networks, Inc.

12. (Optional) Set **Router Lifetime (sec)** to specify how long, in seconds, the client will use the firewall as the default gateway (range is 0-9,000; default is 1,800). Zero specifies that the firewall is not the default gateway. When the lifetime expires, the client removes the firewall entry from its Default Router List and uses another router as the default gateway.
13. Set **Router Preference**, which the client uses to select a preferred router if the network segment has multiple IPv6 routers. **High**, **Medium** (default), or **Low** is the priority that the RA advertises indicating the relative priority of firewall virtual router relative to other routers on the segment.
14. Select **Managed Configuration** to indicate to the client that addresses are available via DHCPv6.
15. Select **Other Configuration** to indicate to the client that other address information (such as DNS-related settings) is available via DHCPv6.
16. Select **Consistency Check** to have the firewall verify that RAs sent from other routers are advertising consistent information on the link. The firewall logs any inconsistencies.
17. Click **OK**.

STEP 7 | (Ethernet or VLAN interface using IPv6 address only) Specify the Recursive DNS Server addresses and DNS Search List the firewall will advertise in ND Router Advertisements from this interface.

The RDNS servers and DNS Search List are part of the DNS configuration for the DNS client so that the client can resolve IPv6 DNS requests.

1. Select **Network** > **Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **IPv6** > **DNS Support**.
4. **Include DNS information in Router Advertisement** to enable the firewall to send IPv6 DNS information.
5. For DNS **Server**, **Add** the IPv6 address of a Recursive DNS Server. **Add** up to eight Recursive DNS servers. The firewall sends server addresses in an ICMPv6 Router Advertisement in order from top to bottom.
6. Specify the **Lifetime** in seconds, which is the maximum length of time the client can use the specific RDNS Server to resolve domain names.
   - The **Lifetime** range is any value equal to or between the **Max Interval** (that you configured on the **Router Advertisement** tab) and two times that **Max Interval**. For example, if your Max Interval is 600 seconds, the Lifetime range is 600-1,200 seconds.
   - The default **Lifetime** is 1,200 seconds.
7. For DNS Suffix, **Add** a **DNS Suffix** (domain name of a maximum of 255 bytes). **Add** up to eight DNS suffixes. The firewall sends suffixes in an ICMPv6 Router Advertisement in order from top to bottom.
8. Specify the **Lifetime** in seconds, which is the maximum length of time the client can use the suffix. The Lifetime has the same range and default value as the **Server**.
9. Click **OK**.

STEP 8 | (Optional) Enable services on the interface.

1. To enable services on the interface, select **Network** > **Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **Advanced** > **Other Info**.
4. Expand the **Management Profile** drop-down, and select a profile or **New Management Profile**.
5. Enter a **Name** for the profile.
6. For **Permitted Services**, select services, such as **Ping**, and click **OK**.

STEP 9 | Commit.

Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

STEP 10 | Cable the interface.

Attach straight through cables from interfaces you configured to the corresponding switch or router on each network segment.

STEP 11 | Verify that the interface is active.

From the web interface, select **Network** > **Interfaces** and verify that icon in the Link State column is green. You can also monitor link state from the **Interfaces** widget on the **Dashboard**.

STEP 12 | Configure static routes and/or a dynamic routing protocol (RIP, OSPF, or BGP) so that the virtual router can route traffic.

- Configure a Static Route
- RIP
- OSPF
- BGP

STEP 13 | Configure a default route.

Configure a Static Route and set it as the default.

Manage IPv6 Hosts Using NDP

This topic describes how you can use NDP to provision IPv6 hosts; therefore, you don't need a separate DHCPv6 server for that purpose. It also explains how to use NDP to monitor IPv6 addresses, allowing you to quickly track the IPv6 address and MAC address of a device and the associated user who has violated a security rule.

- IPv6 Router Advertisements for DNS Configuration
- Configure RDNS Servers and DNS Search List for IPv6 Router Advertisements
- NDP Monitoring
- Enable NDP Monitoring

IPv6 Router Advertisements for DNS Configuration

The firewall implementation of Neighbor Discovery (ND) is enhanced so that you can provision 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 the firewall can provision your 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 IPv6 hosts as part of their DNS configuration to fully provision them to reach internet services. Thus, your IPv6 hosts are configured with:

- The addresses of 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.

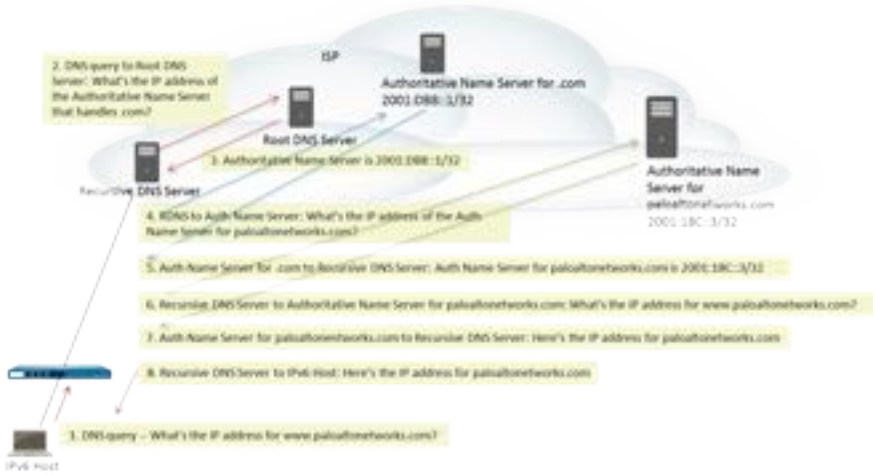
IPv6 Router Advertisement for DNS configuration is supported for Ethernet interfaces, subinterfaces, Aggregated Ethernet interfaces, and Layer 3 VLAN interfaces on all PAN-OS platforms.

The capability of the firewall to send IPv6 RAs for DNS configuration allows the firewall to perform a role similar to DHCP, and is unrelated to the firewall being a DNS proxy, DNS client or DNS server.

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, as shown with

© 2018 Palo Alto Networks, Inc.

three pairs of queries and responses in the following figure. For example, when a user tries to access www.paloaltonetworks.com, the local browser sees that it does not have the IP address for that domain name in its cache, nor does the client's operating system have it. The client's operating system launches a DNS query to a Recursive DNS Server belonging to the local ISP.



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.

A DNS Search List is a list of domain names (suffixes) that the firewall advertises to a DNS client. The firewall thus provisions the DNS client to use the suffixes in its unqualified DNS queries. The DNS client appends the suffixes, one at a time, to an unqualified domain name before it enters the name into a DNS query, thereby using a fully qualified domain name (FQDN) in the DNS query. For example, if a user (of the DNS client being configured) tries to submit a DNS query for the name "quality" without a suffix, the router appends a period and the first DNS suffix from the DNS Search List to the name and transmits a DNS query. If the first DNS suffix on the list is "company.com", the resulting DNS query from the router is for the FQDN "quality.company.com".

If the DNS query fails, the client appends the second DNS suffix from the list to the unqualified name and transmits a new DNS query. The client uses the DNS suffixes in order until a DNS lookup succeeds (ignoring the remaining suffixes) or the router has tried all of the suffixes on the list.

You configure the firewall with the suffixes that you want to provide to the DNS client router in an ND DNSSL option; the DNS client receiving the DNS Search List option is provisioned to use the suffixes in its unqualified DNS queries.

To specify RDNS Servers and a DNS Search List, Configure RDNS Servers and DNS Search List for IPv6 Router Advertisements.

Configure RDNS Servers and DNS Search List for IPv6 Router Advertisements

Perform this task to configure IPv6 Router Advertisements for DNS Configuration of IPv6 hosts.

STEP 1 | Enable the firewall to send IPv6 Router Advertisements from an interface.

1. Select **Network** > **Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface to configure.
3. On the **IPv6** tab, select **Enable IPv6 on the interface**.
4. On the **Router Advertisement** tab, select **Enable Router Advertisement**.
5. Click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 2 | Specify the Recursive DNS Server addresses and DNS Search List the firewall will advertise in ND Router Advertisements from this interface.

The RDNS servers and DNS Search List are part of the DNS configuration for the DNS client so that the client can resolve IPv6 DNS requests.

1. Select **Network** > **Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **IPv6** > **DNS Support**.
4. **Include DNS information in Router Advertisement** to enable the firewall to send IPv6 DNS information.
5. For DNS **Server**, **Add** the IPv6 address of a Recursive DNS Server. **Add** up to eight Recursive DNS servers. The firewall sends server addresses in an ICMPv6 Router Advertisement in order from top to bottom.
6. Specify the **Lifetime** in seconds, which is the maximum length of time the client can use the specific RDNS Server to resolve domain names.
   - The **Lifetime** range is any value equal to or between the **Max Interval** (that you configured on the **Router Advertisement** tab) and two times that **Max Interval**. For example, if your Max Interval is 600 seconds, the Lifetime range is 600-1,200 seconds.
   - The default **Lifetime** is 1,200 seconds.
7. For DNS Suffix, **Add** a **DNS Suffix** (domain name of a maximum of 255 bytes). **Add** up to eight DNS suffixes. The firewall sends suffixes in an ICMPv6 Router Advertisement in order from top to bottom.
8. Specify the **Lifetime** in seconds, which is the maximum length of time the client can use the suffix. The Lifetime has the same range and default value as the **Server**.
9. Click **OK**.

STEP 3 | Commit your changes.

Click **Commit**.

NDP Monitoring

Neighbor Discovery Protocol (NDP) for IPv6 (RFC 4861) performs functions similar to ARP functions for IPv4. The firewall by default runs NDP, which uses ICMPv6 packets to discover and track the link-layer addresses and status of neighbors on connected links.

Enable NDP Monitoring on page 817 so you can view the IPv6 addresses of devices on the link local network, their MAC address, associated username from User-ID (if the user of that device used the directory service to log in), reachability Status of the address, and Last Reported date and time the NDP monitor received a Router Advertisement from this IPv6 address. The username is on a best-case basis; there can be many IPv6 devices on a network with no username, such as printers, fax machines, servers, etc.

If you want to quickly track a device and user who has violated a security rule, it is very useful to have the IPv6 address, MAC address and username displayed all in one place. You need the MAC address that corresponds to the IPv6 address in order to trace the MAC address back to a physical switch or Access Point.



NDP monitoring is not guaranteed to discover all devices because there could be other networking devices between the firewall and the client that filter out NDP or Duplicate Address Detection (DAD) messages. The firewall can monitor only the devices that it learns about on the interface.

NDP monitoring also monitors Duplicate Address Detection (DAD) packets from clients and neighbors. You can also monitor IPv6 ND logs to make troubleshooting easier.

© 2018 Palo Alto Networks, Inc.

NDP monitoring is supported for Ethernet interfaces, subinterfaces, Aggregated Ethernet interfaces, and VLAN interfaces on all PAN-OS models.

Enable NDP Monitoring

Perform this task to enable NDP Monitoring for an interface.

STEP 1 | Enable NDP monitoring.

1. Select **Network** > **Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **IPv6**.
4. Select **Address Resolution**.
5. Select **Enable NDP Monitoring**.

   After you enable or disable NDP monitoring, you must Commit before NDP monitoring can start or stop.
6. Click **OK**.

STEP 2 | Commit your changes.

Click **Commit**.

STEP 3 | Monitor NDP and DAD packets from clients and neighbors.

1. Select **Network** > **Interfaces** and **Ethernet** or **VLAN**.
2. For the interface where you enabled NDP monitoring, in the Features column, hover over the NDP Monitoring icon:

   The NDP Monitoring summary for the interface displays the list of IPv6 **Prefixes** that this interface will send in the Router Advertisement (RA) if RA is enabled (they are the IPv6 prefixes of the interface itself).

   The summary also indicates whether DAD, Router Advertisement, and DNS Support are enabled; IP addresses of any Recursive DNS Servers configured; and any DNS suffixes configured on the DNS Search List.
3. Click on the NDP Monitoring icon to display detailed information.



Each row of the detailed NDP Monitoring table for the interface displays the IPv6 address of a neighbor the firewall has discovered, the corresponding MAC address, corresponding User ID (on a best-case basis), reachability Status of the address, and Last Reported date and time this NDP Monitor received an RA from this IP address. A User ID will not display for printers or other non-

© 2018 Palo Alto Networks, Inc.

user-based hosts. If the status of the IP address is Stale, the neighbor is not known to be reachable, per RFC 4861.

At the bottom right is the count of **Total Devices Detected** on the link local network.

- Enter an IPv6 address in the filter field to search for an address to display.
- Select the check boxes to display or not display IPv6 addresses.
- Click the numbers, the right or left arrow, or the vertical scroll bar to advance through many entries.
- Click **Clear All NDP Entries** to clear the entire table.

STEP 4 | Monitor ND logs for reporting purposes.

1. Select **Monitor** > **Logs** > **System**.
2. In the Type column, view **ipv6nd** logs and corresponding descriptions.

   For example, `inconsistent router advertisement received` indicates that the firewall received an RA different from the RA that it is going to send out.

Configure an Aggregate Interface Group

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 another 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 Link Aggregation Control Protocol (LACP), failure detection is automatic at the physical and data link layers regardless of whether the peers are directly connected. LACP also enables automatic failover to standby interfaces if you configured hot spares. All Palo Alto Networks firewalls except the PA-200 and VM-Series models support aggregate groups. You can add up to eight aggregate groups per firewall and each group can have up to eight interfaces.

PAN-OS firewall models support a maximum of 16,000 IP addresses assigned to physical or virtual Layer 3 interfaces; this maximum includes both IPv4 and IPv6 addresses.

Before configuring an aggregate group, you must configure its interfaces. Among the interfaces assigned to any particular aggregate group, the hardware media can differ (for example, you can mix fiber optic and copper), but the bandwidth and interface type must be the same. The bandwidth and interface type options are:

- **Bandwidth**—1Gbps, 10Gbps, 40Gbps, or 100Gbps
- **Interface type**—HA3, virtual wire, Layer 2, or Layer 3. You can aggregate the HA3 (packet forwarding) interfaces in an active/active high availability (HA) deployment but only for PA-500, PA-3000 Series, and PA-5000 Series firewalls.

This procedure describes configuration steps only for the Palo Alto Networks firewall. You must also configure the aggregate group on the peer device. Refer to the documentation of that device for instructions.

STEP 1 | Configure the general interface group parameters.

1. Select **Network** > **Interfaces** > **Ethernet** and **Add Aggregate Group**.
2. In the field adjacent to the read-only **Interface Name**, enter a number (1–8) to identify the aggregate group.

© 2018 Palo Alto Networks, Inc.

3. For the **Interface Type**, select **HA**, **Virtual Wire**, **Layer2**, or **Layer3**.
4. Configure the remaining parameters for the **Interface Type** you selected.

STEP 2 | Configure the LACP settings.

Perform this step only if you want to enable LACP for the aggregate group.

You cannot enable LACP for virtual wire interfaces.

1. Select the **LACP** tab and **Enable LACP**.
2. Set the **Mode** for LACP status queries to **Passive** (the firewall just responds—the default) or **Active** (the firewall queries peer devices).

   As a best practice, set one LACP peer to active and the other to passive. LACP cannot function if both peers are passive. The firewall cannot detect the mode of its peer device.
3. Set the **Transmission Rate** for LACP query and response exchanges to **Slow** (every 30 seconds—the default) or **Fast** (every second). Base your selection on how much LACP processing your network supports and how quickly LACP peers must detect and resolve interface failures.
4. Select **Fast Failover** if you want to enable failover to a standby interface in less than one second. By default, the option is disabled and the firewall uses the IEEE 802.1ax standard for failover processing, which takes at least three seconds.

   As a best practice, use Fast Failover in deployments where you might lose critical data during the standard failover interval.
5. Enter the **Max Ports** (number of interfaces) that are active (1–8) in the aggregate group. If the number of interfaces you assign to the group exceeds the **Max Ports**, the remaining interfaces will be in standby mode. The firewall uses the **LACP Port Priority** of each interface you assign (Step 3) to determine which interfaces are initially active and to determine the order in which standby interfaces become active upon failover. If the LACP peers have non-matching port priority values, the values of the peer with the lower **System Priority** number (default is 32,768; range is 1–65,535) will override the other peer.
6. (Optional) For active/passive firewalls only, select **Enable in HA Passive State** if you want to enable LACP pre-negotiation for the passive firewall. LACP pre-negotiation enables quicker failover to the passive firewall (for details, see LACP and LLDP Pre-Negotiation for Active/Passive HA).

   If you select this option, you cannot select Same System MAC Address for Active-Passive HA; pre-negotiation requires unique interface MAC addresses on each HA firewall.
7. (Optional) For active/passive firewalls only, select **Same System MAC Address for Active-Passive HA** and specify a single **MAC Address** for both HA firewalls. This option minimizes failover latency if the LACP peers are virtualized (appearing to the network as a single device). By default, the option is disabled: each firewall in an HA pair has a unique MAC address.

   If the LACP peers are not virtualized, use unique MAC addresses to minimize failover latency.

STEP 3 | Assign interfaces to the aggregate group.

Perform the following steps for each interface (1–8) that will be a member of the aggregate group.

1. Select **Network** > **Interfaces** > **Ethernet** and click the interface name to edit it.
2. Set the **Interface Type** to **Aggregate Ethernet**.

© 2018 Palo Alto Networks, Inc.

3. Select the **Aggregate Group** you just defined.
4. Select the **Link Speed**, **Link Duplex**, and **Link State**.

   As a best practice, set the same link speed and duplex values for every interface in the group. For non-matching values, the firewall defaults to the higher speed and full duplex.
5. (Optional) Enter an **LACP Port Priority** (default is 32,768; range is 1–65,535) if you enabled LACP for the aggregate group. If the number of interfaces you assign exceeds the **Max Ports** value of the group, the port priorities determine which interfaces are active or standby. The interfaces with the lower numeric values (higher priorities) will be active.
6. Click **OK**.

STEP 4 | If the firewalls have an active/active configuration and you are aggregating HA3 interfaces, enable packet forwarding for the aggregate group.

1. Select **Device** > **High Availability** > **Active/Active Config** and edit the Packet Forwarding section.
2. Select the aggregate group you configured for the **HA3 Interface** and click **OK**.

STEP 5 | Commit your changes and verify the aggregate group status.

1. Click **Commit**.
2. Select **Network** > **Interfaces** > **Ethernet**.
3. Verify that the Link State column displays a green icon for the aggregate group, indicating that all member interfaces are up. If the icon is yellow, at least one member is down but not all. If the icon is red, all members are down.
4. If you configured LACP, verify that the Features column displays the LACP enabled icon for the aggregate group.

Use Interface Management Profiles to Restrict Access

An Interface Management profile protects the firewall from unauthorized access by defining the protocols, services, and IP addresses that a firewall interface permits for management traffic. For example, you might want to prevent users from accessing the firewall web interface over the ethernet1/1 interface but allow that interface to receive SNMP queries from your network monitoring system. In this case, you would enable SNMP and disable HTTP/HTTPS in an Interface Management profile and assign the profile to ethernet1/1.

You can assign an Interface Management profile to Layer 3 Ethernet interfaces (including subinterfaces) and to logical interfaces (aggregate group, VLAN, loopback, and tunnel interfaces). If you do not assign an Interface Management profile to an interface, it denies access for all IP addresses, protocols, and services by default.

The management (MGT) interface does not require an Interface Management profile. You restrict protocols, services, and IP addresses for the MGT interface when you Perform Initial Configuration of the firewall. In case the MGT interface goes down, allowing management access over another interface enables you to continue managing the firewall. When enabling access to a firewall interface using an Interface Management profile, make sure you do not enable management access (HTTP, HTTPS, SSH, or Telnet) from the internet or from other untrusted zones inside your enterprise security boundary. Follow the Best Practices for Securing Administrative Access to ensure that you are properly securing management access to your firewall.

STEP 1 | Configure the Interface Management profile.

© 2018 Palo Alto Networks, Inc.

1. Select **Network** > **Network Profiles** > **Interface Mgmt** and click **Add**.
2. Select the protocols that the interface permits for management traffic: **Ping**, **Telnet**, **SSH**, **HTTP**, **HTTP OCSP**, **HTTPS**, or **SNMP**.
3. Select the services that the interface permits for management traffic:
   - **Response Pages**—Use to enable response pages for:
     - **Captive Portal**—To serve Captive Portal response pages, the firewall leaves ports open on Layer 3 interfaces: port 6080 for NT LAN Manager (NTLM), 6081 for Captive Portal in transparent mode, and 6082 for Captive Portal in redirect mode. For details, see Configure Captive Portal.
     - **URL Admin Override**—For details, see Allow Password Access to Certain Sites.
   - **User-ID**—Use to Redistribute User Mappings and Authentication Timestamps.
   - **User-ID Syslog Listener-SSL** or **User-ID Syslog Listener-UDP**—Use to Configure User-ID to Monitor Syslog Senders for User Mapping over SSL or UDP.
4. (Optional) **Add** the Permitted IP Addresses that can access the interface. If you don't add entries to the list, the interface has no IP address restrictions.
5. Click **OK**.

STEP 2 | Assign the Interface Management profile to an interface.

1. Select **Network** > **Interfaces**, select the type of interface (**Ethernet**, **VLAN**, **Loopback**, or **Tunnel**), and select the interface.
2. Select **Advanced** > **Other info** and select the Interface **Management Profile** you just added.
3. Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Virtual Routers

The firewall uses virtual routers to obtain routes to other subnets by you manually defining 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 firewall's IP routing information base (RIB). 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 forwarding information base (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 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, 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, OSPF, OSPFv3, or RIP) as well as add static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that aren't shared between virtual routers, enabling 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. While each interface can belong to only one virtual router, you can configure multiple routing protocols and static routes for a virtual router. Regardless of the static routes and dynamic routing protocols you configure for a virtual router, one general configuration is required:

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

- Interfaces on the firewall that you want to perform routing.
- Administrative distances for static, OSPF internal, OSPF external, IBGP, EBGP and RIP.

STEP 2 | Create a virtual router and apply interfaces to it.

The firewall comes with a virtual router named **default**. You can edit the **default** virtual router or add a new virtual router.

1. Select **Network** > **Virtual Routers**.
2. Select a virtual router (the one named **default** or a different virtual router) or **Add** the **Name** of a new virtual router.
3. Select **Router Settings** > **General**.
4. Click **Add** in the **Interfaces** box and select an already defined interface from the drop-down.

   Repeat this step for all interfaces you want to add to the virtual router.
5. Click **OK**.

STEP 3 | Set Administrative Distances for static and dynamic routing.

Set Administrative Distances for types of routes as required for your network. When the virtual router has two or more different routes to the same destination, it uses administrative distance to choose the best path from different routing protocols and static routes, by preferring a lower distance.

- **Static**—Range is 10-240; default is 10.
- **OSPF Internal**—Range is 10-240; default is 30.
- **OSPF External**—Range is 10-240; default is 110.

© 2018 Palo Alto Networks, Inc.

- **IBGP**—Range is 10-240; default is 200.
- **EBGP**—Range is 10-240; default is 20.
- **RIP**—Range is 10-240; default is 120.



See ECMP if you want to leverage having multiple equal-cost paths for forwarding.

STEP 4 | Commit virtual router general settings.

Click **OK** and **Commit**.

STEP 5 | Configure Ethernet, VLAN, loopback, and tunnel interfaces as needed.

Configure Layer 3 Interfaces.

© 2018 Palo Alto Networks, Inc.

Service Routes

The firewall uses the management (MGT) interface by default to access external services, such as DNS servers, external authentication servers, 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 (shown in the following task) or Customize Service Routes for a Virtual System on a firewall enabled for multiple virtual systems so that you have the flexibility to use interfaces associated with a virtual system. 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.

The following procedure enables you to change the interface the firewall uses to send requests to external services.

STEP 1 | Customize service routes.

1. Select **Device** > **Setup** > **Services** > **Global** (omit Global on a firewall without multiple virtual system capability), and in the Services Features section, click **Service Route Configuration**.



2. Select **Customize** and do one of the following to create a service route:
   - For a predefined service:
     - Select **IPv4** or **IPv6** and click the link for the service for which you want customize the service route.

       To easily use the same source address for multiple services, select the checkbox for the services, click Set Selected Routes, and proceed to the next step.
     - To limit the drop-down list for Source Address, select a **Source Interface**; then select a **Source Address** (from that interface) as the service route. Selecting **Any** Source Interface makes all IP addresses on all interfaces available in the Source Address drop-down from which you select an address. Selecting **Use default** causes the firewall to use the management interface for the service route, unless the packet destination IP address matches the configured Destination IP address, in which case the source IP address is set to the **Source Address** configured for the **Destination**. Selecting **MGT** causes the firewall to use the MGT interface for the service route, regardless of any destination service route.
     - Click **OK** to save the setting.
     - Repeat this step if you want to specify both an IPv4 and IPv6 address for a service.
   - For a destination service route:
     - Select **Destination** and **Add** a **Destination** IP address. In this case, if a packet arrives with a destination IP address that matches this configured **Destination** address, then the source IP address of the packet will be set to the **Source Address** configured in the next step.
     - To limit the drop-down list for Source Address, select a **Source Interface**; then select a **Source Address** (from that interface) as the service route. Selecting **Any** Source Interface makes all IP addresses on all interfaces available in the Source Address drop-down from which you select an address. Selecting **MGT** causes the firewall to use the MGT interface for the service route.
     - Click **OK** to save the setting.
3. Repeat the prior steps for each service route you want to customize.

© 2018 Palo Alto Networks, Inc.

4. Click **OK** to save the service route configuration.

STEP 2 | Commit.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Static Routes

Static routes are typically used in conjunction with dynamic routing protocols. You might configure a static route for a location that a dynamic routing protocol can't reach. Static routes require manual configuration on every router in the network, rather than the firewall entering dynamic routes in its route tables; even though static routes require that configuration on all routers, they may be desirable in small networks rather than configuring a routing protocol.

- Static Route Overview
- Static Route Removal Based on Path Monitoring
- Configure a Static Route
- Configure Path Monitoring for a Static Route

Static Route Overview

If you decide that you want specific Layer 3 traffic to take a certain route without participating in IP routing protocols, you can Configure a Static Route on page 828 using IPv4 and IPv6 routes.

A default route is a specific static route. If you don't use dynamic routing to obtain a default route for your virtual router, you must configure a static default route. When the virtual router has an incoming packet and finds no match for the packet's destination in its route table, the virtual router sends the packet to the default route. The default IPv4 route is 0.0.0.0/0; the default IPv6 route is ::/0. You can configure both an IPv4 and IPv6 default route

Static routes themselves don't change or adjust to changes in network environments, so traffic typically isn't rerouted if a failure occurs along the route to a statically defined endpoint. However, you have options to back up static routes in the event of a problem:

- You can configure a static route with a Bidirectional Forwarding Detection (BFD on page 948) profile so that if a BFD session between the firewall and the BFD peer fails, the firewall removes the failed static route from the RIB and FIB tables and uses an alternate path with a lower priority.
- You can Configure Path Monitoring for a Static Route on page 829 so that the firewall can use an alternative route.

By default, static routes have an administrative distance of 10. When the firewall has two or more routes to the same destination, it uses the route with the lowest administrative distance. By increasing the administrative distance of a static route to a value higher than a dynamic route, you can use the static route as a backup route if the dynamic route is unavailable.

While you're configuring a static route, you can specify whether the firewall installs an IPv4 static route in the unicast or multicast route table (RIB), or both tables, or doesn't install the route at all. For example, you could install an IPv4 static route in the multicast route table only, because you want only multicast traffic to use that route. This option give you more control over which route the traffic takes. You can specify whether the firewall installs an IPv6 static route in the unicast route table or not.

Static Route Removal Based on Path Monitoring

When you Configure Path Monitoring for a Static Route, the firewall uses path monitoring to detect when the path to one or more monitored destination has gone down. The firewall can then reroute traffic using an alternative routes. The firewall uses path monitoring for static routes much like path monitoring for HA or policy-based forwarding (PBF), as follows:

- The firewall sends ICMP ping messages (heartbeat messages) to one or more monitored destinations that you determine are robust and reflect the availability of the static route.

© 2018 Palo Alto Networks, Inc.

- If pings to any or all of the monitored destinations fail, the firewall considers the static route down too and removes it from the Routing Information Base (RIB) and Forwarding Information Base (FIB). The RIB is the table of static routes the firewall is configured with and dynamic routes it has learned from routing protocols. The FIB is the forwarding table of routes the firewall uses for forwarding packets. The firewall selects an alternative static route to the same destination (based on the route with the lowest metric) from the RIB and places it in the FIB.
- The firewall continues to monitor the failed route. When the route comes back up, and (based on the **Any** or **All** failure condition) the path monitor returns to Up state, the preemptive hold timer begins. The path monitor must remain up for the duration of the hold timer; then the firewall considers the static route stable and reinstates it into the RIB. The firewall then compares metrics of routes to the same destination to decide which route goes in the FIB.

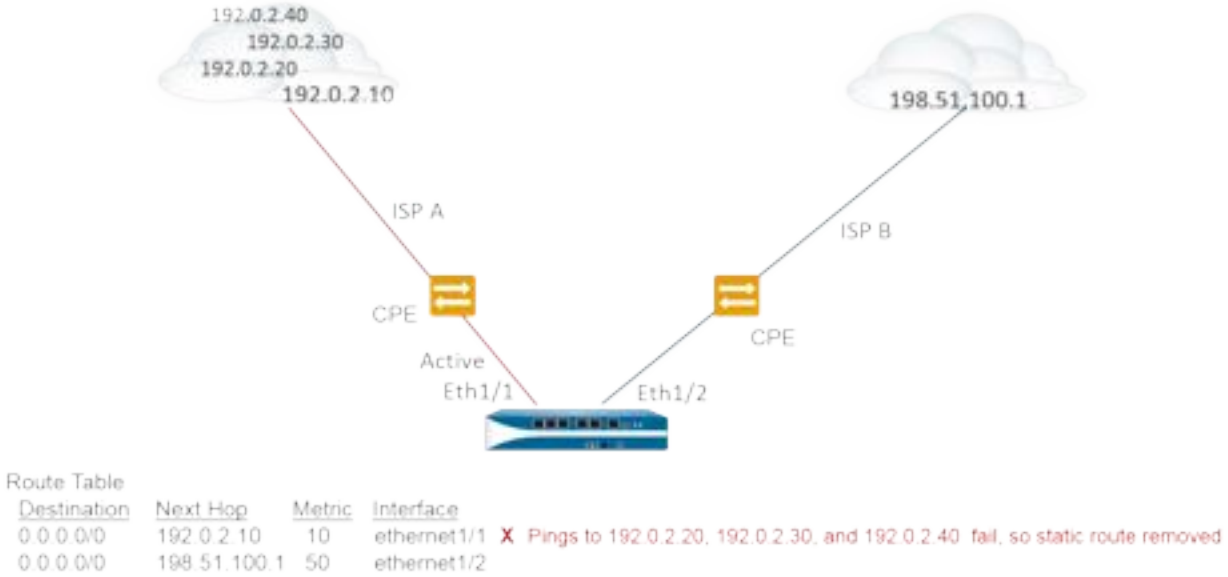
Path monitoring is a desirable mechanism to avoid blackholing traffic for:

- A static or default route.
- A static or default route redistributed into a routing protocol.
- A static or default route when one peer does not support BFD. (The best practice is not to enable both BFD and path monitoring on a single interface.)
- A static or default route instead of using PBF path monitoring, which doesn't remove a failed static route from the RIB, FIB, or redistribution policy.

Path monitoring doesn't apply to static routes configured between virtual routers.

In the following figure, the firewall is connected to two ISPs for route redundancy to the internet. The primary default route 0.0.0.0 (metric 10) uses Next Hop 192.0.2.10; the secondary default route 0.0.0.0 (metric 50) uses Next Hop 198.51.100.1. The customer premises equipment (CPE) for ISP A keeps the primary physical link active, even after internet connectivity goes down. With the link artificially active, the firewall can't detect that the link is down and that it should replace the failed route with the secondary route in its RIB.

To avoid blackholing traffic to a failed link, configure path monitoring of 192.0.2.20, 192.0.2.30, and 192.0.2.40 and if all (or any) of the paths to these destinations fail, the firewall presumes the path to Next Hop 192.0.2.10 is also down, removes the static route 0.0.0.0 (that uses Next Hop 192.0.2.10) from its RIB, and replaces it with the secondary route to the same destination 0.0.0.0 (that uses Next Hop 198.51.100.1), which also accesses the internet.



© 2018 Palo Alto Networks, Inc.

When you Configure a Static Route, one of the required fields is the Next Hop toward that destination. The type of next hop you configure determines the action the firewall takes during path monitoring, as follows:

If Next Hop Type in Static Route is:	Firewall Action for ICMP Ping
IP Address	The firewall uses the source IP address and egress interface of the static route as the source address and egress interface in the ICMP ping. It uses the configured Destination IP address of the monitored destination as the ping's destination address. It uses the static route's next hop address as the ping's next hop address.
Next VR	The firewall uses the source IP address of the static route as the source address in the ICMP ping. The egress interface is based on the lookup result from the next hop's virtual router. The configured Destination IP address of the monitored destination is the ping's destination address.
None	The firewall uses the destination IP address of the path monitor as the next hop and sends the ICMP ping to the interface specified in the static route.

When path monitoring for a static or default route fails, the firewall logs a critical event (path-monitor-failure). When the static or default route recovers, the firewall logs another critical event (path-monitor-recovery).

Firewalls synchronize path monitoring configurations for an active/passive HA deployment, but the firewall blocks egress ICMP ping packets on a passive HA peer because it is not actively processing traffic. The firewall doesn't synchronize path monitoring configurations for active/active HA deployments.

Configure a Static Route

Perform the following task to configure Static Routes or a default route for a virtual router on the firewall.

STEP 1 | Configure a static route.

1. Select **Network** > **Virtual Router** and select the virtual router you are configuring, such as **default**.
2. Select the **Static Routes** tab.
3. Select **IPv4** or **IPv6**, depending on the type of static route you want to configure.
4. **Add** a **Name** for the route.
5. 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).
6. (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.
7. 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).
   - **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.
   - **Discard**—Select to drop packets that are addressed to this destination.

© 2018 Palo Alto Networks, Inc.

- **None**—Select if there is no next hop for the route. For example, a point-to-point connection needs no next hop because there is only one way for packets to go.

8. Enter an **Admin Distance** for the route to override the default administrative distance set for static routes for this virtual router (range is 10-240; default is 10).
9. Enter a **Metric** for the route (range is 1-65,535).

STEP 2 | Choose where to install the route.

Select the **Route Table** (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 used only for unicast traffic.
- **Multicast**—Install the route in the multicast route table (available for IPv4 routes only). Choose this option if you want the route used only for multicast traffic.
- **Both**—Install the route in the unicast and multicast route tables (available for 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 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. Default is **None**.

STEP 4 | Click **OK** twice.

STEP 5 | Commit the configuration.

Click **Commit**.

Configure Path Monitoring for a Static Route

Use the following procedure to configure Static Route Removal Based on Path Monitoring.

STEP 1 | Enable path monitoring for a static route.

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **Static Routes**, select **IPv4** or **IPv6**, and select the static route you want to monitor. You can monitor up to 128 static routes.
3. Select **Path Monitoring** to enable path monitoring for the route.

STEP 2 | Configure the monitored destination(s) for the static route.

1. **Add** a monitored destination by **Name**. You can add up to eight monitored destinations per static route.
2. Select **Enable** to monitor the destination.
3. For **Source IP**, select the IP address that the firewall uses in the ICMP ping to the monitored destination:
   - If the interface has multiple IP addresses, select one.
   - If you select an interface, the firewall uses the first IP address assigned to the interface by default.
   - If you select **DHCP (Use DHCP Client address)**, the firewall uses the address that DHCP assigned to the interface. To see the DHCP address, select **Network** > **Interfaces** > **Ethernet** and in the row for the Ethernet interface, click on **Dynamic DHCP Client**. The IP Address displays in the Dynamic IP Interface Status window.
4. For **Destination IP**, enter an IP address or address object to which the firewall will monitor the path. The monitored destination and static route destination must use the same address family (IPv4 or IPv6).

© 2018 Palo Alto Networks, Inc.



The destination IP address should belong to a reliable endpoint; you wouldn't want to base path monitoring on a device that itself is unstable or unreliable.

5. (Optional) Specify the ICMP **Ping Interval (sec)** in seconds to determine how frequently the firewall monitors the path (range is 1-60; default is 3).
6. (Optional) Specify the ICMP **Ping Count** of packets that don't return from the destination before the firewall considers the static route down and removes it from the RIB and FIB (range is 3-10; default is 5).
7. Click **OK**.

STEP 3 | Determine whether path monitoring for the static route is based on one or all monitored destinations, and set the preemptive hold time.

1. Select a **Failure Condition**, whether **Any** or **All** of the monitored destinations for the static route must be unreachable by ICMP for the firewall to remove the static route from the RIB and FIB and add the static route that has the next lowest metric going to the same destination to the FIB.



Select All to avoid the possibility of any single monitored destination signaling a route failure when the destination is simply offline for maintenance, for example.

2. (Optional) Specify the **Preemptive Hold Time (min)**, which is the number of minutes a downed path monitor must remain in Up state before the firewall reinstalls the static route into the RIB. The path monitor evaluates all of its monitored destinations for the static route and comes up based on the **Any** or **All** failure condition. If a link goes down or flaps during the hold time, when the link comes back up, the path monitor can come back up; the timer restarts when the path monitor returns to Up state.

   A **Preemptive Hold Time** of zero causes the firewall to reinstall the route into the RIB immediately upon the path monitor coming up. Range is 0-1,440; default is 2.
3. Click **OK**.

STEP 4 | Commit.

Click **Commit**.

STEP 5 | Verify path monitoring on static routes.

1. Select **Network** > **Virtual Routers** and in the row of the virtual router you are interested in, select **More Runtime Stats**.
2. From the **Routing** tab, select **Static Route Monitoring**.
3. For a static route (Destination), view whether Path Monitoring is Enabled or Disabled. The Status column indicates whether the route is Up, Down, or Disabled. Flags for the static route are: A—active, S—static, E—ECMP.
4. Select **Refresh** periodically to see the latest state of the path monitoring (health check).
5. Hover over the Status of a route to view the monitored IP addresses and results of the pings sent to the monitored destinations for that route. For example, 3/5 means that a ping interval of 3 seconds and a ping count of 5 consecutive missed pings (the firewall receives no ping in the last 15 seconds) indicates path monitoring detects a link failure. Based on the **Any** or **All** failure condition, if path monitoring is in failed state and the firewall receives a ping after 15 seconds, the path can be deemed up and the **Preemptive Hold Time** starts.

   The State indicates the last monitored ping results: success or failed. Failed indicates that the series of ping packets (ping interval multiplied by ping count) was not successful. A single ping packet failure does not reflect a failed ping state.

STEP 6 | View the RIB and FIB to verify that the static route is removed.

© 2018 Palo Alto Networks, Inc.

1. Select **Network** > **Virtual Routers** and in the row of the virtual router you are interested in, select **More Runtime Stats**.
2. From the **Routing** tab, select **Route Table** (RIB) and then the **Forwarding Table** (FIB) to view each, respectively.
3. Select **Unicast** or **Multicast** to view the appropriate route table.
4. For **Display Address Family**, select **IPv4 and IPv6**, **IPv4 Only**, or **IPv6 Only**.
5. (Optional) In the filter field, enter the route you are searching for and select the arrow, or use the scroll bar to move through pages of routes.
6. See if the route is removed or present.
7. Select **Refresh** periodically to see the latest state of the path monitoring (health check).



To view the events logged for path monitoring, select Monitor > Logs > System. View the entry for path-monitor-failure, which indicates path monitoring for a static route destination failed, so the route was removed. View the entry for path-monitor-recovery, which indicates path monitoring for the static route destination recovered, so the route was restored.

© 2018 Palo Alto Networks, Inc.

RIP

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

Perform the following procedure to configure RIP.

STEP 1 | Configure general virtual router configuration settings.

See Virtual Routers for details.

STEP 2 | Configure general RIP configuration settings.

1. Select the **RIP** tab.
2. Select **Enable** to enable the RIP protocol.
3. Select **Reject Default Route** if you do not want to learn any default routes through RIP. This is the recommended, default setting.

   Clear **Reject Default Route** if you want to permit redistribution of default routes through RIP.

STEP 3 | Configure interfaces for RIP.

1. On the **Interfaces** tab, select an interface from the drop-down in the Interface configuration section.
2. Select an already defined interface.
3. Select **Enable**.
4. Select **Advertise** to advertise a default route to RIP peers with the specified metric value.
5. (Optional) Select a profile from the **Auth Profile** drop-down.
6. Select normal, passive or send-only from the **Mode** drop-down.
7. Click **OK**.

STEP 4 | Configure RIP timers.

1. On the **Timers** tab, enter a value for **Interval Seconds (sec)**. This setting defines the length of the following RIP timer intervals in seconds (range is 1-60; default is 1).
2. Specify the **Update Intervals** to define the number of intervals between route update announcements (range is 1-3,600; default is 30).
3. Specify the **Delete Intervals** to define the number of intervals between the time that the route expires to its deletion (range is 1-3,600; default is 180).
4. Specify the **Expire Intervals** to define the number of intervals between the time that the route was last updated to its expiration (range is 1-3600; default is 120).

STEP 5 | (Optional) Configure Auth Profiles.

By default, the firewall does not use RIP authentication for the exchange between RIP neighbors. Optionally, you can configure RIP authentication between RIP neighbors by either a simple password or MD5 authentication. MD5 authentication is recommended; it is more secure than a simple password.

Simple Password RIP authentication

1. Select **Auth Profiles** and **Add** a name for the authentication profile to authenticate RIP messages.
2. Select **Simple Password** as the **Password Type**.
3. Enter a simple password and then confirm.

© 2018 Palo Alto Networks, Inc.

MD5 RIP authentication

1. Select **Auth Profiles** and **Add** a name for the authentication profile to authenticate RIP messages.
2. Select **MD5** as the **Password Type**.
3. **Add** one or more password entries, including:
   - Key-ID (range is 0-255)
   - Key
4. (Optional) Select **Preferred** status.
5. Click **OK** to specify the key to be used to authenticate outgoing message.
6. Click **OK** again in the Virtual Router - RIP Auth Profile dialog box.

STEP 6 | **Commit** your changes.

© 2018 Palo Alto Networks, Inc.

OSPF

Open Shortest Path First (OSPF) is an interior gateway protocol (IGP) that is most often used to dynamically manage network routes in large enterprise network. It determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The information gathered from the LSAs is used to construct a topology map of the network. This topology map is shared across routers in the network and used to populate the IP routing table with available routes.

Changes in the network topology are detected dynamically and used to generate a new topology map within seconds. A shortest path tree is computed of each route. Metrics associated with each routing interface are used to calculate the best route. These can include distance, network throughput, link availability etc. Additionally, these metrics can be configured statically to direct the outcome of the OSPF topology map.

The Palo Alto Networks implementation of OSPF fully supports the following RFCs:

- RFC 2328 (for IPv4)
- RFC 5340 (for IPv6)

The following topics provide more information about the OSPF and procedures for configuring OSPF on the firewall:

- OSPF Concepts
- Configure OSPF
- Configure OSPFv3
- Configure OSPF Graceful Restart
- Confirm OSPF Operation

OSPF Concepts

The following topics introduce the OSPF concepts you must understand in order to configure the firewall to participate in an OSPF network:

- OSPFv3
- OSPF Neighbors
- OSPF Areas
- OSPF Router Types

OSPFv3

OSPFv3 provides support for the OSPF routing protocol within an IPv6 network. As such, it provides support for IPv6 addresses and prefixes. It retains most of the structure and functions in OSPFv2 (for IPv4) with some minor changes. The following are some of the additions and changes to OSPFv3:

- **Support for multiple instances per link**—With OSPFv3, you can run multiple instances of the OSPF protocol over a single link. This is accomplished by assigning an OSPFv3 instance ID number. An interface that is assigned to an instance ID drops packets that contain a different ID.
- **Protocol Processing Per-link**—OSPFv3 operates per-link instead of per-IP-subnet as on OSPFv2.
- **Changes to Addressing**—IPv6 addresses are not present in OSPFv3 packets, except for LSA payloads within link state update packets. Neighboring routers are identified by the Router ID.
- **Authentication Changes**—OSPFv3 doesn't include any authentication capabilities. Configuring OSPFv3 on a firewall requires an authentication profile that specifies Encapsulating Security Payload (ESP) or IPv6 Authentication Header (AH).The re-keying procedure specified in RFC 4552 is not supported in this release.

© 2018 Palo Alto Networks, Inc.

- **Support for multiple instances per-link**—Each instance corresponds to an instance ID contained in the OSPFv3 packet header.
- **New LSA Types**—OSPFv3 supports two new LSA types: Link LSA and Intra Area Prefix LSA.

All additional changes are described in detail in RFC 5340.

OSPF Neighbors

Two OSPF-enabled routers connected by a common network and in the same OSPF area that form a relationship are OSPF neighbors. The connection between these routers can be through a common broadcast domain or by a point-to-point connection. This connection is made through the exchange of hello OSPF protocol packets. These neighbor relationships are used to exchange routing updates between routers.

OSPF Areas

OSPF operates within a single autonomous system (AS). Networks within this single AS, however, can be divided into a number of areas. By default, Area 0 is created. Area 0 can either function alone or act as the OSPF backbone for a larger number of areas. Each OSPF area is named using a 32-bit identifier which in most cases is written in the same dotted-decimal notation as an IP4 address. For example, Area 0 is usually written as 0.0.0.0.

The topology of an area is maintained in its own link state database and is hidden from other areas, which reduces the amount of traffic routing required by OSPF. The topology is then shared in a summarized form between areas by a connecting router.

OSPF Area Type	Description
Backbone Area	The backbone area (Area 0) is the core of an OSPF network. All other areas are connected to it and all traffic between areas must traverse it. All routing between areas is distributed through the backbone area. While all other OSPF areas must connect to the backbone area, this connection doesn't need to be direct and can be made through a virtual link.
Normal OSPF Area	In a normal OSPF area there are no restrictions; the area can carry all types of routes.
Stub OSPF Area	A stub area does not receive routes from other autonomous systems. Routing from the stub area is performed through the default route to the backbone area.
NSSA Area	The Not So Stubby Area (NSSA) is a type of stub area that can import external routes, with some limited exceptions.

OSPF Router Types

Within an OSPF area, routers are divided into the following categories.

- **Internal Router**—A router with that has OSPF neighbor relationships only with devices in the same area.
- **Area Border Router (ABR)**—A router that has OSPF neighbor relationships with devices in multiple OSPF areas. ABRs gather topology information from their connected areas and distribute it to the backbone area.
- **Backbone Router**—A backbone router is a router that runs OSPF and has at least one interface connected to the OSPF backbone area. Since ABRs are always connected to the backbone, they are always classified as backbone routers.

© 2018 Palo Alto Networks, Inc.

- **Autonomous System Boundary Router (ASBR)**—An ASBR is a router that attaches to more than one routing protocol and exchanges routing information between them.

Configure OSPF

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

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

STEP 1 | Configure general virtual router configuration settings.

See Virtual Routers for details.

STEP 2 | Enable OSPF.

1. Select the **OSPF** tab.
2. Select **Enable** to enable the OSPF protocol.
3. Enter the **Router ID**.
4. Select **Reject Default Route** if you do not want to learn any default routes through OSPF. This is the recommended, default setting.

   Clear **Reject Default Route** if you want to permit redistribution of default routes through OSPF.

STEP 3 | Configure Areas - Type for the OSPF protocol.

1. On the **Areas** tab, **Add** an **Area ID** for the area in *x.x.x.x* format. This is the identifier that each neighbor must accept to be part of the same area.
2. On the **Type** tab, select one of the following from the area **Type** drop-down:
   - **Normal**—There are no restrictions; the area can carry all types of routes.
   - **Stub**—There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, configure the following:
     - **Accept Summary**—Link state advertisements (LSA) are accepted from other areas. If this option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.
     - **Advertise Default Route**—Default route LSAs will be included in advertisements to the stub area along with a configured metric value in the configured range 1-255.
   - **NSSA** (Not-So-Stubby Area)—The firewall can leave the area only by routes other than OSPF routes. If you select NSSA, select **Accept Summary** and **Advertise Default Route** as described for **Stub**. If you select this option, configure the following:
     - **Type**—Select either **Ext 1** or **Ext 2** route type to advertise the default LSA.
     - **Ext Ranges**—**Add** ranges of external routes that you want to **Advertise** or for which you want to **Suppress** advertising.
3. Click **OK**.

STEP 4 | Configure Areas - Range for the OSPF protocol

1. On the **Range** tab, **Add** aggregate LSA destination addresses in the area into subnets.

© 2018 Palo Alto Networks, Inc.

2. **Advertise** or **Suppress** advertising LSAs that match the subnet, and click **OK**. Repeat to add additional ranges.

STEP 5 | Configure Areas - Interfaces for the OSPF protocol

1. On the **Interface** tab, **Add** the following information for each interface to be included in the area:
   - **Interface**—Select an interface from the drop-down.
   - **Enable**—Selecting this option causes the OSPF interface settings to take effect.
   - **Passive**—Select if you do not want the OSPF interface to send or receive OSPF packets. Although OSPF packets are not sent or received if you choose this option, the interface is included in the LSA database.
   - **Link type**—Choose **Broadcast** if you want all neighbors that are accessible through the interface to be discovered automatically by multicasting OSPF hello messages, such as an Ethernet interface. Choose **p2p** (point-to-point) to automatically discover the neighbor. Choose **p2mp** (point-to-multipoint) when neighbors must be defined manually and **Add** the neighbor IP addresses for all neighbors that are reachable through this interface.
   - **Metric**—Enter an OSPF metric for this interface (range is 0-65,535; default is 10).
   - **Priority**—Enter an OSPF priority for this interface. This is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR) (range is 0-255; default is 1). If zero is configured, the router will not be elected as a DR or BDR.
   - **Auth Profile**—Select a previously-defined authentication profile.
   - **Timing**—Modify the timing settings if desired (not recommended). For details on these settings, refer to the online help.
2. Click **OK**.

STEP 6 | Configure Areas - Virtual Links.

1. On the **Virtual Link** tab, **Add** the following information for each virtual link to be included in the backbone area:
   - **Name**—Enter a name for the virtual link.
   - **Enable**—Select to enable the virtual link.
   - **Neighbor ID**—Enter the router ID of the router (neighbor) on the other side of the virtual link.
   - **Transit Area**—Enter the area ID of the transit area that physically contains the virtual link.
   - **Timing**—It is recommended that you keep the default timing settings.
   - **Auth Profile**—Select a previously-defined authentication profile.
2. Click **OK** to save virtual links.
3. Click **OK** to save area.

STEP 7 | (Optional) Configure Auth Profiles.

By default, the firewall does not use OSPF authentication for the exchange between OSPF neighbors. Optionally, you can configure OSPF authentication between OSPF neighbors by either a simple password or using MD5 authentication. MD5 authentication is recommended; it is more secure than a simple password.

Simple Password OSPF authentication

1. Select the **Auth Profiles** tab and **Add** a name for the authentication profile to authenticate OSPF messages.
2. Select **Simple Password** as the **Password Type**.
3. Enter a simple password and then confirm.

MD5 OSPF authentication

1. Select the **Auth Profiles** tab and **Add** a name for the authentication profile to authenticate OSPF messages.

© 2018 Palo Alto Networks, Inc.

2. Select **MD5** as the **Password Type** and **Add** one or more password entries, including:
   - Key-ID (range is 0-255)
   - Key
   - Select the **Preferred** option to specify that the key be used to authenticate outgoing messages.
3. Click **OK**.

STEP 8 | Configure Advanced OSPF options.

1. On the **Advanced** tab, select **RFC 1583 Compatibility** to ensure compatibility with RFC 1583.
2. Specify a value for the **SPF Calculation Delay (sec)** timer, which allows you to tune the delay time (in seconds) between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should use the same delay value to optimize convergence times.
3. Specify a value for the **LSA Interval (sec)** timer, which is the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur.
4. Click **OK**.

STEP 9 | **Commit** your changes.

Configure OSPFv3

OSPF supports both IPv4 and IPv6. You must use OSPFv3 if you are using IPv6.

STEP 1 | Configure general virtual router configuration settings.

See Virtual Routers for details.

STEP 2 | Configure general OSPFv3 configuration settings.

1. Select the **OSPFv3** tab.
2. Select **Enable** to enable the OSPF protocol.
3. Enter the **Router ID**.
4. Select **Reject Default Route** if you do not want to learn any default routes through OSPFv3 This is the recommended default setting.

   Clear **Reject Default Route** if you want to permit redistribution of default routes through OSPFv3.

STEP 3 | Configure Auth Profile for the OSPFv3 protocol.

While OSPFv3 doesn't include any authentication capabilities of its own, it relies entirely on IPSec to secure communications between neighbors.

When configuring an authentication profile, you must use Encapsulating Security Payload (ESP) (recommended) or IPv6 Authentication Header (AH).

ESP OSPFv3 authentication

1. On the **Auth Profiles** tab, **Add** a name for the authentication profile to authenticate OSPFv3 messages.
2. Specify a Security Policy Index (**SPI**) (hexadecimal value in the range from 00000000 to FFFFFFFF). The two ends of the OSPFv3 adjacency must have matching SPI values.
3. Select **ESP** for **Protocol**.
4. Select a **Crypto Algorithm** from the drop-down.

   You can select **None** or one of the following algorithms: **SHA1**, **SHA256**, **SHA384**, **SHA512**, or **MD5**.
5. If a **Crypto Algorithm** other than None was selected, enter a value for **Key** and then confirm.

© 2018 Palo Alto Networks, Inc.

AH OSPFv3 authentication

1. On the **Auth Profiles** tab, **Add** a name for the authentication profile to authenticate OSPFv3 messages.
2. Specify a Security Policy Index (**SPI**). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a hexadecimal value between 00000000 and FFFFFFFF.
3. Select **AH** for **Protocol**.
4. Select a **Crypto Algorithm** from the drop-down.

   You must enter one of the following algorithms: **SHA1**, **SHA256**, **SHA384**, **SHA512**, or **MD5**.
5. Enter a value for **Key** and then confirm.
6. Click **OK**.
7. Click **OK** again in the Virtual Router - OSPF Auth Profile dialog.

STEP 4 | Configure Areas - Type for the OSPFv3 protocol.

1. On the **Areas** tab, **Add** an **Area ID**. This is the identifier that each neighbor must accept to be part of the same area.
2. On the **General** tab, select one of the following from the area **Type** drop-down:
   - **Normal**—There are no restrictions; the area can carry all types of routes.
   - **Stub**—There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, configure the following:
     - **Accept Summary**—Link state advertisements (LSA) are accepted from other areas. If this option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.
     - **Advertise Default Route**—Default route LSAs will be included in advertisements to the stub area along with a configured metric value in the configured range 1-255.
   - **NSSA** (Not-So-Stubby Area)—The firewall can leave the area only by routes other than OSPF routes. If selected, configure **Accept Summary** and **Advertise Default Route** as described for **Stub**. If you select this option, configure the following:
     - **Type**—Select either **Ext 1** or **Ext 2** route type to advertise the default LSA.
     - **Ext Ranges**—**Add** ranges of external routes that you want to enable or suppress advertising for.

STEP 5 | Associate an OSPFv3 authentication profile to an area or an interface.

To an Area

1. On the **Areas** tab, select an existing area from the table.
2. On the **General** tab, select a previously defined **Authentication Profile** from the **Authentication** drop-down.
3. Click **OK**.

To an Interface

1. On the **Areas** tab, select an existing area from the table.
2. Select the **Interface** tab and **Add** the authentication profile you want to associate with the OSPF interface from the **Auth Profile** drop-down.
3. Click **OK**.

STEP 6 | Click **OK** again to save the area settings.

STEP 7 | (Optional) Configure Export Rules.

© 2018 Palo Alto Networks, Inc.

1. On the **Export Rules** tab, select **Allow Redistribute Default Route** to permit redistribution of default routes through OSPFv3.
2. Click **Add**.
3. Enter the **Name**; the value must be a valid IPv6 subnet or valid redistribution profile name.
4. Select **New Path Type**, **Ext 1** or **Ext 2**.
5. Specify a **New Tag** for the matched route, using has a 32-bit value in dotted-decimal notation.
6. Assign a **Metric** to the new rule (range is 1-16,777,215).
7. Click **OK**.

STEP 8 | Configure Advanced OSPFv3 options.

1. On the **Advanced** tab, select **Disable Transit Routing for SPF Calculation** if you want the firewall to participate in OSPF topology distribution without being used to forward transit traffic.
2. Specify a value for the **SPF Calculation Delay (sec)** timer, which allows you to tune the delay time (in seconds) between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should use the same delay value to optimize convergence times.
3. Specify a value for the **LSA Interval (sec)** timer, which is the minimum time (in seconds) between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur.
4. (Optional) Configure OSPF Graceful Restart.
5. Click **OK**.

STEP 9 | **Commit** your changes.

Configure OSPF Graceful Restart

OSPF Graceful Restart directs OSPF neighbors to continue using routes through a firewall during a short transition when it is out of service. This behavior increases network stability by reducing the frequency of routing table reconfiguration and the related route flapping that can occur during short periodic down times.

For a Palo Alto Networks firewall, OSPF Graceful Restart involves the following operations:

- **Firewall as a restarting device**—If the firewall will be down for a short period of time or is unavailable for short intervals, it sends Grace LSAs to its OSPF neighbors. The neighbors must be configured to run in Graceful Restart helper mode. In helper mode, the neighbor receives Grace LSAs informing it that the firewall will perform a graceful restart within a specified period of time defined as the **Grace Period**. During the grace period, the neighbor continues to forward routes through the firewall and to send LSAs that announce routes through the firewall. If the firewall resumes operation before expiration of the grace period, traffic forwarding will continue as before without network disruption. If the firewall does not resume operation after the grace period has expired, the neighbors will exit helper mode and resume normal operation, which will involve reconfiguring the routing table to bypass the firewall.
- **Firewall as a Graceful Restart Helper**—If neighboring routers may be down for short periods of time, the firewall can be configured to operate in Graceful Restart helper mode, in which case the firewall employs a **Max Neighbor Restart Time**. When the firewall receives the Grace LSAs from its OSPF neighbor, it continues to route traffic to the neighbor and advertise routes through the neighbor until either the grace period or max neighbor restart time expires. If neither expires before the neighbor returns to service, traffic forwarding continues as before without network disruption. If either period expires before the neighbor returns to service, the firewall exits helper mode and resumes normal operation, which involves reconfiguring the routing table to bypass the neighbor.

© 2018 Palo Alto Networks, Inc.

STEP 1 | Select **Network** > **Virtual Routers** and select the virtual router you want to configure.

STEP 2 | Select **OSPF** > **Advanced** or **OSPFv3** > **Advanced**.

STEP 3 | Verify that the following are selected (they are enabled by default):

- **Enable Graceful Restart**
- **Enable Helper Mode**
- **Enable Strict LSA Checking**

These should remain selected unless required by your topology.

STEP 4 | Configure a **Grace Period** in seconds.

STEP 5 | Configure a **Max Neighbor Restart Time** in seconds.

Confirm OSPF Operation

Once an OSPF configuration has been committed, you can use any of the following operations to confirm that OSPF is operating:

- View the Routing Table on page 841
- Confirm OSPF Adjacencies on page 841
- Confirm that OSPF Connections are Established on page 841

View the Routing Table

By viewing the routing table, you can see whether OSPF routes have been established. The routing table is accessible from either the web interface or the CLI. If you are using the CLI, use the following commands:

- `show routing route`
- `show routing fib`

If you are using the web interface to view the routing table, use the following workflow:

STEP 1 | Select **Network** > **Virtual Routers** and in the same row as the virtual router you are interested in, click the **More Runtime Stats** link.

STEP 2 | Select **Routing** > **Route Table** and examine the **Flags** column of the routing table for routes that were learned by OSPF.

Confirm OSPF Adjacencies

Use the following workflow to confirm that OSPF adjacencies have been established:

STEP 1 | Select **Network** > **Virtual Routers** and in the same row as the virtual router you are interested in, click the **More Runtime Stats** link.

STEP 2 | Select **OSPF** > **Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established.

Confirm that OSPF Connections are Established

View the System log to confirm that the firewall has established OSPF connections.

© 2018 Palo Alto Networks, Inc.

STEP 1 | Select **Monitor** > **System** and look for messages to confirm that OSPF adjacencies have been established.

STEP 2 | Select **OSPF** > **Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established (are full).

© 2018 Palo Alto Networks, Inc.

BGP

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

- BGP Overview on page 843
- MP-BGP on page 843
- Configure BGP on page 845
- Configure a BGP Peer with MP-BGP for IPv4 or IPv6 Unicast on page 849
- Configure a BGP Peer with MP-BGP for IPv4 Multicast on page 851

BGP Overview

BGP functions between autonomous systems (exterior BGP or eBGP) or within an AS (interior BGP or iBGP) to exchange routing and reachability information with BGP speakers. The firewall provides a complete BGP implementation, which includes the following features:

- Specification of one BGP routing instance per virtual router.
- BGP settings per virtual router, which include basic parameters such as local route ID and local AS, and advanced options such as path selection, route reflector, BGP Confederations, route flap dampening, and graceful restart.
- Peer group and neighbor settings, which include neighbor address and remote AS, and advanced options such as neighbor attributes and connections.
- Route policies to control route import, export and advertisement; prefix-based filtering; and address aggregation.
- IGP-BGP interaction to inject routes to BGP using redistribution profiles.
- Authentication profiles, which specify the MD5 authentication key for BGP connections. Authentication helps prevent route leaking and successful DoS attacks.
- Multiprotocol BGP (MP-BGP) to allow BGP peers to carry IPv6 unicast routes and IPv4 multicast routes in Update packets, and to allow the firewall and a BGP peer to communicate with each other using IPv6 addresses.

MP-BGP

BGP supports IPv4 unicast prefixes, but a BGP network that uses IPv4 multicast routes or IPv6 unicast prefixes needs multiprotocol BGP (MP-BGP) in order to exchange routes of address types other than IPv4 unicast. MP-BGP allows BGP peers to carry IPv4 multicast routes and IPv6 unicast routes in Update packets, in addition to the IPv4 unicast routes that BGP peers can carry without MP-BGP enabled.

In this way, MP-BGP provides IPv6 connectivity to your BGP networks that use either native IPv6 or dual stack IPv4 and IPv6. Service providers can offer IPv6 service to their customers, and enterprises can use IPv6 service from service providers. The firewall and a BGP peer can communicate with each other using IPv6 addresses.

In order for BGP to support multiple network-layer protocols (other than BGP for IPv4), Multiprotocol Extensions for BGP-4 (RFC 4760) use Network Layer Reachability Information (NLRI) in a Multiprotocol Reachable NLRI attribute that the firewall sends and receives in BGP Update packets. That attribute contains information about the destination prefix, including these two identifiers:

- The Address Family Identifier (AFI), as defined by the IANA in Address Family Numbers, indicates that the destination prefix is an IPv4 or IPv6 address. (PAN-OS supports IPv4 and IPv6 AFIs.)

© 2018 Palo Alto Networks, Inc.

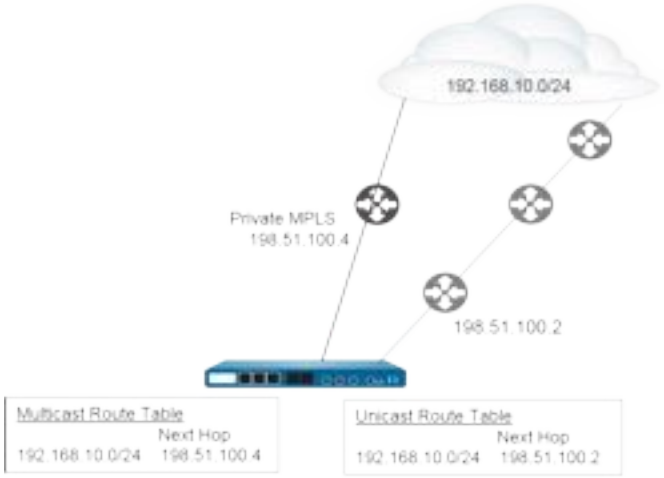
- The Subsequent Address Family Identifier (SAFI) in PAN-OS indicates that the destination prefix is a unicast or multicast address (if the AFI is IPv4), or that the destination prefix is a unicast address (if the AFI is IPv6). PAN-OS does not support IPv6 multicast.

If you enable MP-BGP for IPv4 multicast or if you configure a multicast static route, the firewall supports separate unicast and multicast route tables for static routes. You might want to separate the unicast and multicast traffic going to the same destination. The multicast traffic can take a different path from unicast traffic because, for example, your multicast traffic is critical, so you need it to be more efficient by having it take fewer hops or undergo less latency.

You can also exercise more control over how BGP functions by configuring BGP to use routes from only the unicast or multicast route table (or both) when BGP imports or exports routes, sends conditional advertisements, or performs route redistribution or route aggregation.

You can decide to use a dedicated multicast RIB (route table) by enabling MP-BGP and selecting the Address Family of IPv4 and Subsequent Address Family of multicast or by installing an IPv4 static route in the multicast route table. After you do either of those methods to use the multicast RIB, the firewall uses the multicast RIB for all multicast routing and reverse path forwarding (RPF). If you prefer to use the unicast RIB for all routing (unicast and multicast), you should not enable the multicast RIB by either method.

In the following figure, a static route to 192.168.10.0/24 is installed in the unicast route table, and its next hop is 198.51.100.2. However, multicast traffic can take a different path to a private MPLS cloud; the same static route is installed in the multicast route table with a different next hop (198.51.100.4) so that its path is different.



Using separate unicast and multicast route tables gives you more flexibility and control when you configure these BGP functions:

- Install an IPv4 static route into the unicast or multicast route table, or both, as described in the preceding example. (You can install an IPv6 static route into the unicast route table only).
- Create an Import rule so that any prefixes that match the criteria are imported into the unicast or multicast route table, or both.
- Create an Export rule so that prefixes that match the criteria are exported (sent to a peer) from the unicast or multicast route table, or both.
- Configure a conditional advertisement with a Non Exist filter so that the firewall searches the unicast or multicast route table (or both) to ensure the route doesn't exist in that table, and so the firewall advertises a different route.
- Configure a conditional advertisement with an Advertise filter so that the firewall advertises routes matching the criteria from the unicast or multicast route table, or both.
- Redistribute a route that appears in the unicast or multicast route table, or both.

© 2018 Palo Alto Networks, Inc.

- Configure route aggregation with an advertise filter so that aggregated routes to be advertised come from the unicast or multicast route table, or both.
- Conversely, configure route aggregation with a suppress filter so that aggregated routes that should be suppressed (not advertised) come from the unicast or multicast route table, or both.

When you configure a peer with MP-BGP using an Address Family of IPv6, you can use IPv6 addresses in the Address Prefix and Next Hop fields of an Import rule, Export rule, Conditional Advertisement (Advertise Filter and Non Exist Filter), and Aggregate rule (Advertise Filter, Suppress Filter, and Aggregate Route Attribute).

Configure BGP

Perform the following task to configure BGP.

STEP 1 | Configure general virtual router configuration settings.

See Virtual Routers for details.

STEP 2 | Enable BGP for the virtual router, assign a router ID, and assign the virtual router to an AS.

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **BGP**.
3. Select **Enable** to enable BGP for this virtual router.
4. Assign a **Router ID** to BGP for the virtual router, which is typically an IPv4 address to ensure the Router ID is unique.
5. Assign the **AS Number**, the number of the AS to which the virtual router belongs, based on the router ID. Range is 1-4,294,967,295.
6. Click **OK**.

STEP 3 | Configure general BGP configuration settings.

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **BGP** > **General**.
3. Select **Reject Default Route** to ignore any default routes that are advertised by BGP peers.
4. Select **Install Route** to install BGP routes in the global routing table.
5. Select **Aggregate MED** to enable route aggregation even when routes have different Multi-Exit Discriminator (MED) values.
6. Specify the **Default Local Preference** that can be used to determine preferences among different paths.
7. Select the **AS Format** for interoperability purposes:
   - **2 Byte** (default value)
   - **4 Byte**
8. Enable or disable each of the following settings for **Path Selection**:
   - **Always Compare MED**—Enable this comparison to choose paths from neighbors in different autonomous systems.
   - **Deterministic MED Comparison**—Enable this comparison to choose between routes that are advertised by IBGP peers (BGP peers in the same autonomous system).
9. For **Auth Profiles**, **Add** an authentication profile:
   - **Profile Name**—Enter a name to identify the profile.
   - **Secret/Confirm Secret**—Enter and confirm a passphrase for BGP peer communications. The Secret is used as a key in MD5 authentication.
10. Click **OK**.
11. Click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 4 | (Optional) Configure BGP settings.

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **BGP** > **Advanced**.
3. Select **ECMP Multiple AS Support** if you configured ECMP and you want to run ECMP over multiple BGP autonomous systems.
4. Select **Enforce First AS for EBGP** to cause the firewall to drop an incoming Update packet from an eBGP peer that doesn't list the eBGP peer's own AS number as the first AS number in the AS_PATH attribute. Default is enabled.
5. Select **Graceful Restart** and configure the following timers:
   - **Stale Route Time (sec)**—Specifies the length of time in seconds that a route can stay in the stale state (range is 1-3,600; default is 120).
   - **Local Restart Time (sec)**—Specifies the length of time in seconds that the local device waits to restart. This value is advertised to peers (range is 1-3,600; default is 120).
   - **Max Peer Restart Time (sec)**—Specifies the maximum length of time in seconds that the local device accepts as a grave period restart time for peer devices (range is 1-3,600; default is 120).
6. For **Reflector Cluster ID**, specify an IPv4 identifier to represent the reflector cluster.
7. For **Confederation Member AS**, specify the autonomous system number identifier (also called a sub-AS number), which is visible only within the BGP confederation. For more information, see BGP Confederations.
8. **Add** the following information for each Dampening Profile that you want to configure, select **Enable**, and click **OK**:
   - **Profile Name**—Enter a name to identify the profile.
   - **Cutoff**—Specify a route withdrawal threshold above which a route advertisement is suppressed (range is 0.0-1,000.0; default is 1.25).
   - **Reuse**—Specify a route withdrawal threshold below which a suppressed route is used again (range is 0.0-1,000.0; default is 5).
   - **Max Hold Time (sec)**—Specify the maximum length of time in seconds that a route can be suppressed, regardless of how unstable it has been (range is 0-3,600 seconds; default is 900).
   - **Decay Half Life Reachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered reachable (range is 0-3,600 seconds; default is 300).
   - **Decay Half Life Unreachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered unreachable (range is 0-3,600; default is 300).
9. Click **OK**.
10. Click **OK**.

STEP 5 | Configure a BGP peer group.

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **BGP** > **Peer Group** and **Add** a **Name** for the peer group and select **Enable**.
3. Select **Aggregated Confed AS Path** to include a path to the configured aggregated confederation AS.
4. Select **Soft Reset with Stored Info** to perform a soft reset of the firewall after updating the peer settings.
5. Select the **Type** of peer group:
   - **IBGP**—**Export Next Hop**: Select **Original** or **Use self**
   - **EBGP Confed**—**Export Next Hop**: Select **Original** or **Use self**
   - **EBGP Confed**—**Export Next Hop**: Select **Original** or **Use self**
   - **EBGP**—**Import Next Hop**: Select **Original** or **Use self**, **Export Next Hop**: Specify **Resolve** or **Use self**. Select **Remove Private AS** if you want to force BGP to remove private AS numbers from the AS_PATH attribute in Updates that the firewall sends to a peer in another AS.
6. Click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 6 | Configure a BGP peer that belongs to the peer group and specify its addressing.

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **BGP** > **Peer Group** and select the peer group you created.
3. For Peer, **Add** a peer by **Name**.
4. Select **Enable** to activate the peer.
5. Enter the **Peer AS** to which the peer belongs.
6. Select **Addressing**.
7. For **Local Address**, select the **Interface** for which you are configuring BGP. If the interface has more than one **IP** address, enter the IP address for that interface to be the BGP peer.
8. For **Peer Address**, enter the **IP** address of the BGP peer.
9. Click **OK**.

STEP 7 | Configure connection settings for the BGP peer.

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **BGP** > **Peer Group** and select the peer group you created.
3. Select the **Peer** you configured.
4. Select **Connection Options**.
5. Select an **Auth Profile** for the peer.
6. Set a **Keep Alive Interval (sec)**, the interval (in seconds) after which routes from the peer are suppressed according to the Hold Time setting (range is 0-1,200; default is 30).
7. Set **Multi Hop**, the time-to-live (TTL) value in the IP header (range is 1-255; default is 0. The default value of 0 means 2 for eBGP prior to PAN-OS 8.0.2, and it means 1 beginning with PAN-OS 8.0.2. The default value of 0 means 255 for iBGP).
8. Set **Open Delay Time (sec)**, the delay in seconds between a TCP handshake and the firewall sending the first BGP Open message to establish a BGP connection (range is 0-240; default is 0).
9. Set **Hold Time (sec)**, the length of time in seconds that may elapse between successive Keepalive or Update messages from the peer before the peer connection is closed (range is 3-3,600; default is 90).
10. Set **Idle Hold Time (sec)**, the length of time to wait (in seconds) before retrying to connect to the peer (range is 1-3,600; default is 15).
11. For **Incoming Connections**, enter a **Remote Port** and select **Allow** to allow incoming traffic to this port.
12. For **Outgoing Connections**, enter a **Local Port** and select **Allow** to allow outgoing traffic from this port.
13. Click **OK**.

STEP 8 | Configure the BGP peer with settings for route reflector client, peering type, maximum prefixes, and Bidirectional Forwarding Detection (BFD).

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **BGP** > **Peer Group** and select the peer group you created.
3. Select the **Peer** you configured.
4. Select **Advanced**.
5. For **Reflector Client**, select one of the following:
   - **non-client**—Peer is not a route reflector client (default).
   - **client**—Peer is a route reflector client.
   - **meshed-client**
6. For **Peering Type**, select one of the following:
   - **Bilateral**—The two BGP peers establish a peer connection.
   - **Unspecified**— (default).

© 2018 Palo Alto Networks, Inc.

7. For **Max Prefixes**, enter the maximum number of supported IP prefixes (range is 1-100,000) or select **unlimited**.
8. To enable **BFD** for the peer (and thereby override the BFD setting for BGP, as long as BFD is not disabled for BGP at the virtual router level), select one of the following:
   - **default**—Peer uses only default BFD settings.
   - **Inherit-vr-global-setting** (default)—Peer inherits the BFD profile that you selected globally for BGP for the virtual router.
   - A BFD profile you configured—See Create a BFD profile.

   Selecting Disable BFD disables BFD for the BGP peer.
9. Click **OK**.

STEP 9 | Configure Import and Export rules.

The import/export rules are used to import/export routes from/to other routers. For example, importing the default route from your Internet Service Provider.

1. Select the **Import** tab and **Add** a name in the **Rules** field and select **Enable**.
2. **Add** the **Peer Group** from which the routes will be imported.
3. Click the **Match** tab and define the options used to filter routing information. You can also define the Multi-Exit Discriminator (MED) value and a next hop value to routers or subnets for route filtering. The MED option is an external metric that lets neighbors know about the preferred path into an AS. A lower value is preferred over a higher value.
4. Click the **Action** tab and define the action that should occur (allow/deny) based on the filtering options defined in the **Match** tab. If **Deny** is selected, no further options need to be defined. If the **Allow** action is selected, define the other attributes.
5. Click the **Export** tab and define export attributes, which are similar to the **Import** settings, but are used to control route information that is exported from the firewall to neighbors.
6. Click **OK**.

STEP 10 | Configure conditional advertising, which allows you to control what route to advertise in the event that a different route is not available in the local BGP routing table (LocRIB), indicating a peering or reachability failure.

This is useful in cases where you want to try to force routes to one AS over another, for example if you have links to the Internet through multiple ISPs and you want traffic to be routed to one provider instead of the other unless there is a loss of connectivity to the preferred provider.

1. Select the **Conditional Adv** tab, **Add** a name in the **Policy** field.
2. Select **Enable**.
3. **Add** in the **Used By** section the peer group(s) that will use the conditional advertisement policy.
4. Select the **Non Exist Filter** tab and define the network prefix(es) of the preferred route. This specifies the route that you want to advertise, if it is available in the local BGP routing table. If a prefix is going to be advertised and matches a Non Exist filter, the advertisement will be suppressed.
5. Select the **Advertise Filters** tab and define the prefix(es) of the route in the Local-RIB routing table that should be advertised in the event that the route in the non-exist filter is not available in the local routing table. If a prefix is going to be advertised and does not match a Non Exist filter, the advertisement will occur.

STEP 11 | Configure aggregate options to summarize routes in the BGP configuration.

BGP route aggregation is used to control how BGP aggregates addresses. Each entry in the table results in one aggregate address being created. This will result in an aggregate entry in the routing table when at least one or more specific route matching the address specified is learned.

© 2018 Palo Alto Networks, Inc.

1. Select the **Aggregate** tab, and **Add** a name for the aggregate address.
2. In the **Prefix** field, enter the network prefix that will be the primary prefix for the aggregated prefixes.
3. Select the **Suppress Filters** tab and define the attributes that will cause the matched routes to be suppressed.
4. Select the **Advertise Filters** tab and define the attributes that will cause the matched routes to always be advertised to peers.

STEP 12 | Configure redistribution rules.

This rule is used to redistribute host routes and unknown routes that are not on the local RIB to the peers routers.

1. Select the **Redist Rules** tab and click **Add**.
2. In the **Name** field, enter an IP subnet or select a redistribution profile. You can also configure a new redistribution profile from the drop-down if needed.
3. **Enable** the rule.
4. In the **Metric** field, enter the route metric that will be used for the rule.
5. In the **Set Origin** drop-down, select **incomplete**, **igp**, or **egp**.
6. (Optional) Set MED, local preference, AS path limit and community values.

Configure a BGP Peer with MP-BGP for IPv4 or IPv6 Unicast

After you Configure BGP, configure a BGP peer with MP-BGP for IPv4 or IPv6 unicast for either of the following reasons:

- To have your BGP peer carry IPv6 unicast routes, configure MP-BGP with the Address Family Type of **IPv6** and Subsequent Address Family of **Unicast** so that the peer can send BGP updates that include IPv6 unicast routes. BGP peering (Local Address and Peer Address) can still both be IPv4 addresses, or they can both be IPv6 addresses.
- To perform BGP peering over IPv6 addresses (**Local Address** and **Peer Address** use IPv6 addresses).

The following task shows how to enable a BGP peer with MP-BGP so it can carry IPv6 unicast routes, and so it can peer using IPv6 addresses.

The task also shows how to view the unicast or multicast route tables, and how to view the forwarding table, the BGP local RIB, and BGP RIB Out (routes sent to neighbors) to see routes from the unicast or multicast route table or a specific address family (IPv4 or IPv6).

STEP 1 | Enable MP-BGP Extensions for a peer.

Configure the following so that a BGP peer can carry IPv4 or IPv6 unicast routes in Updates packets and the firewall can use IPv4 or IPv6 addresses to communicate with its peer.

1. Select **Network** > **Virtual Routers** and select the virtual router you are configuring.
2. Select **BGP**.
3. Select **Peer Group** and select a peer group.
4. Select a BGP peer (router).
5. Select **Addressing**.
6. Select **Enable MP-BGP Extensions** for the peer.
7. For **Address Family Type**, select **IPv4** or **IPv6**. For example, select IPv6.
8. For **Subsequent Address Family**, **Unicast** is selected. If you chose **IPv4** for the Address Family, you can select **Multicast** also.
9. For **Local Address**, select an **Interface** and optionally select an **IP** address, for example, 2001:DB8:55::/32

© 2018 Palo Alto Networks, Inc.

10. For **Peer Address**, enter the peer's **IP** address, using the same address family (IPv4 or IPv6) as the Local Address, for example, 2001:DB8:58::/32.
11. Select **Advanced**.
12. (Optional) **Enable Sender Side Loop Detection**. When you enable Sender Side Loop Detection, the firewall will check the AS_PATH attribute of a route in its FIB before it sends the route in an update, to ensure that the peer AS number is not on the AS_PATH list. If it is, the firewall removes it to prevent a loop
13. Click **OK**.

STEP 2 | (Optional) Create a static route and install it in the unicast route table because you want the route to be used only for unicast purposes.

1. Select **Network** > **Virtual Routers** and select the virtual router you are configuring.
2. Select **Static Routes**, select **IPv4** or **IPv6**, and **Add** a route.
3. Enter a **Name** for the static route.
4. Enter the IPv4 or IPv6 **Destination** prefix and netmask, depending on whether you chose IPv4 or IPv6.
5. Select the egress **Interface**.
6. Select the **Next Hop** as **IPv6 Address** (or **IP Address** if you chose IPv4) and enter the address of the next hop to which you want to direct unicast traffic for this static route.
7. Enter an **Admin Distance**.
8. Enter a **Metric**.
9. For **Route Table**, select **Unicast**.
10. Click **OK**.

STEP 3 | Commit the configuration.

Click **Commit**.

STEP 4 | View the unicast or multicast route table.

1. Select **Network** > **Virtual Routers**.
2. In the row for the virtual router, click **More Runtime Stats**.
3. Select **Routing** > **Route Table**.
4. For **Route Table**, select **Unicast** or **Multicast** to display only those routes.
5. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.



Selecting Multicast with IPv6 Only is not supported.

STEP 5 | View the Forwarding Table.

1. Select **Network** > **Virtual Routers**.
2. In the row for the virtual router, click **More Runtime Stats**.
3. Select **Routing** > **Forwarding Table**.
4. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.

STEP 6 | View the BGP RIB tables.

1. View the BGP Local RIB, which shows the BGP routes that the firewall uses to route BGP packets.
    1. Select **Network** > **Virtual Routers**.
    2. In the row for the virtual router, click **More Runtime Stats**.

© 2018 Palo Alto Networks, Inc.

3. Select **BGP** > **Local RIB**.
4. For **Route Table**, select **Unicast** or **Multicast** to display only those routes.
5. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.



Selecting Multicast with IPv6 Only is not supported.

2. View the BGP RIB Out table, which shows the routes that the firewall sends to BGP neighbors.
   1. Select **Network** > **Virtual Routers**.
   2. In the row for the virtual router, click **More Runtime Stats**.
   3. Select **BGP** > **RIB Out**.
   4. For **Route Table**, select **Unicast** or **Multicast** to display only those routes.
   5. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.



Selecting Multicast with IPv6 Only is not supported.

Configure a BGP Peer with MP-BGP for IPv4 Multicast

After you Configure BGP, configure a BGP peer with MP-BGP for IPv4 multicast if you want your BGP peer to be able to learn and pass IPv4 multicast routes in BGP updates. You'll be able to separate unicast from multicast traffic, or employ the features listed in MP-BGP to use only routes from the unicast or multicast route table, or routes from both tables.

If you want to support multicast traffic only, you must use a filter to eliminate unicast traffic.

The firewall doesn't support ECMP for multicast traffic.

STEP 1 | Enable MP-BGP extensions so that a BGP peer can exchange IPv4 multicast routes.

1. Select **Network** > **Virtual Routers** and select the virtual router you are configuring.
2. Select **BGP**.
3. Select **Peer Group**, select a peer group and a BGP peer.
4. Select **Addressing**.
5. Select **Enable MP-BGP Extensions**.
6. For **Address Family Type**, select **IPv4**.
7. For **Subsequent Address Family**, select **Unicast** and then **Multicast**.
8. Click **OK**.

STEP 2 | (Optional) Create an IPv4 static route and install it in the multicast route table only.

You would do this to direct multicast traffic for a BGP peer to a specific next hop, as shown in the topology in MP-BGP.

1. Select **Network** > **Virtual Routers** and select the virtual router you are configuring.
2. Select **Static Routes** > **IPv4** and **Add** a **Name** for the route.
3. Enter the IPv4 **Destination** prefix and netmask.
4. Select the egress **Interface**.
5. Select the **Next Hop** as **IP Address** and enter the IP address of the next hop to which you want to direct multicast traffic for this static route.
6. Enter an **Admin Distance**.
7. Enter a **Metric**.
8. For **Route Table**, select **Multicast**.

© 2018 Palo Alto Networks, Inc.

9. Click **OK**.

STEP 3 | Commit the configuration.

Click **Commit**.

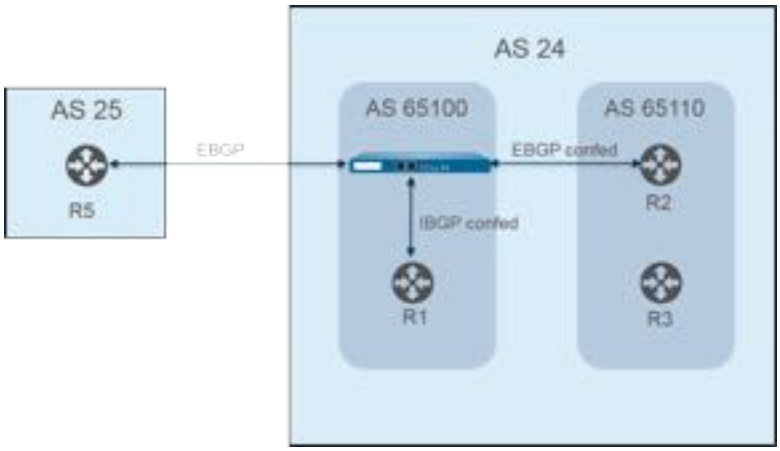
STEP 4 | View the route table.

1. Select **Network** > **Virtual Routers**.
2. In the row for the virtual router, click **More Runtime Stats**.
3. Select **Routing** > **Route Table**.
4. For **Route Table**, select **Unicast** or **Multicast** to display only those routes.
5. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.

STEP 5 | To view the Forwarding table, BGP Local RIB, or BGP RIB Out table, see Configure a BGP Peer with MP-BGP for IPv4 or IPv6 Unicast.

BGP Confederations

BGP confederations provide a way to divide an autonomous system (AS) into two or more sub-autonomous systems (sub-AS) to reduce the burden that the full mesh requirement for IBGP causes. The firewalls (or other routing devices) within a sub-AS must still have a full iBGP mesh with the other firewalls in the same sub-AS. You need BGP peering between sub-autonomous systems for full connectivity within the main AS. The firewalls peering with each other within a sub-AS form an IBGP confederation peering. The firewall in one sub-AS peering with a firewall in a different sub-AS form an EBGP confederation peering. Two firewalls from different autonomous systems that connect are EBGP peers.



Autonomous systems are identified with a public (globally-assigned) AS number, such as AS 24 and AS 25 in the preceding figure. In a PAN-OS environment, you assign each sub-AS a unique Confederation Member AS number, which is a private number seen only within the AS. In this figure, the confederations are AS 65100 and AS 65110. (RFC6996, Autonomous System (AS) Reservation for Private Use, indicates that the IANA reserves AS numbers 64512-65534 for private use.)

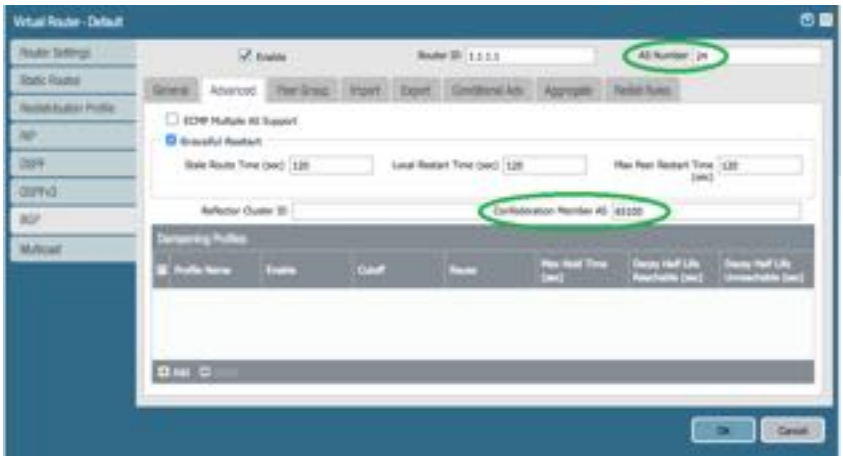
The sub-AS confederations seems like full autonomous systems to each other within the AS. However, when the firewall sends an AS path to an EBGP peer, only the public AS number appears in the AS path; no private sub-AS (Confederation Member AS) numbers are included.

BGP peering occurs between the firewall and R2; the firewall in the figure has these relevant configuration settings:

- AS number—24

© 2018 Palo Alto Networks, Inc.

- Confederation Member AS—65100
- Peering Type—EBGP confed
- Peer AS—65110



Router 2 (R2) in AS 65110 is configured as follows:

- AS number—24
- Confederation Member AS—65110
- Peering Type—EBGP confed
- Peer AS—65100

BGP peering also occurs between the firewall and R1. The firewall has the following additional configuration:

- AS number—24
- Confederation Member AS—65100
- Peering Type—IBGP confed
- Peer AS—65110

R1 is configured as follows:

- AS number—24
- Confederation Member AS—65110
- Peering Type—IBGP confed
- Peer AS—65100

BGP peering occurs between the firewall and R5. The firewall has the following additional configuration:

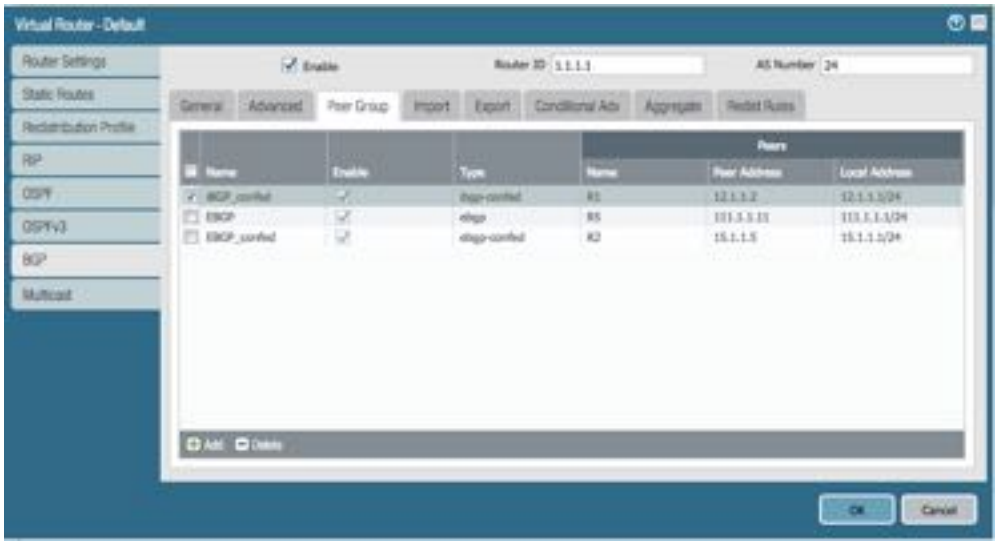
- AS number—24
- Confederation Member AS—65100
- Peering Type—EBGP
- Peer AS—25

R5 is configured as follows:

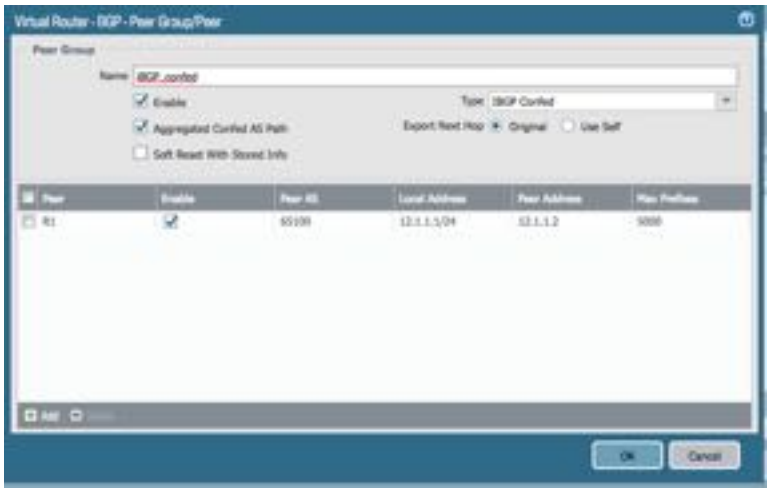
- AS—25
- Peering Type—EBGP
- Peer AS—24

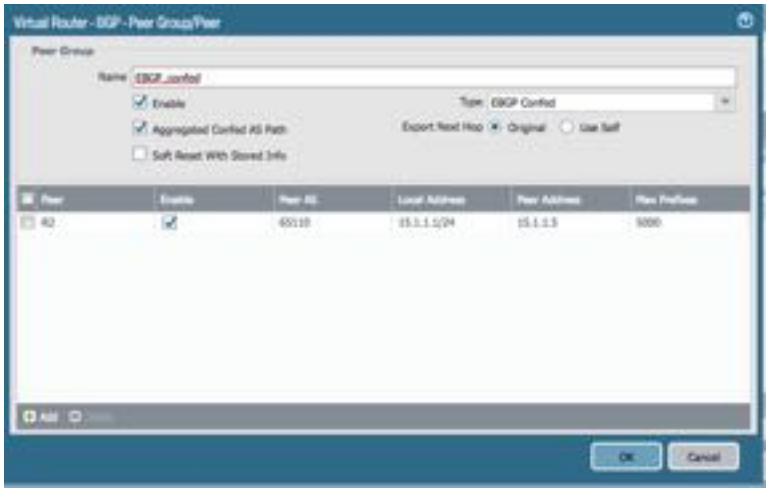
After the firewall is configured to peer with R1, R2, and R5, its peers are visible on the **Peer Group** tab:

© 2018 Palo Alto Networks, Inc.

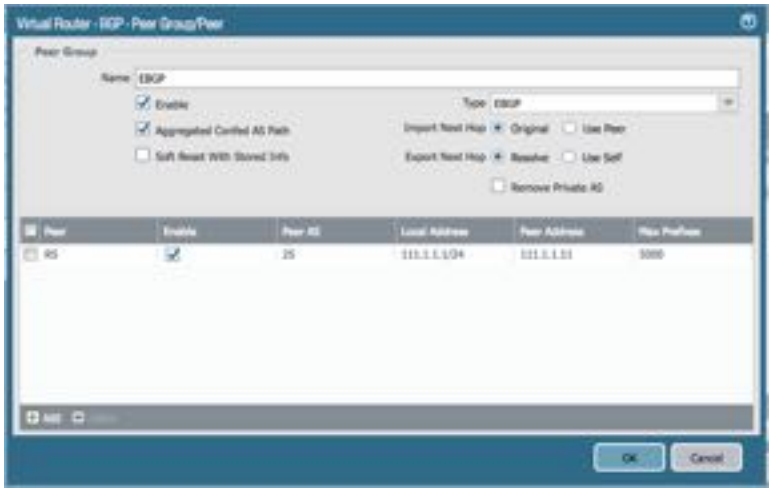


The firewall shows the R1, R2, and R5 peers:

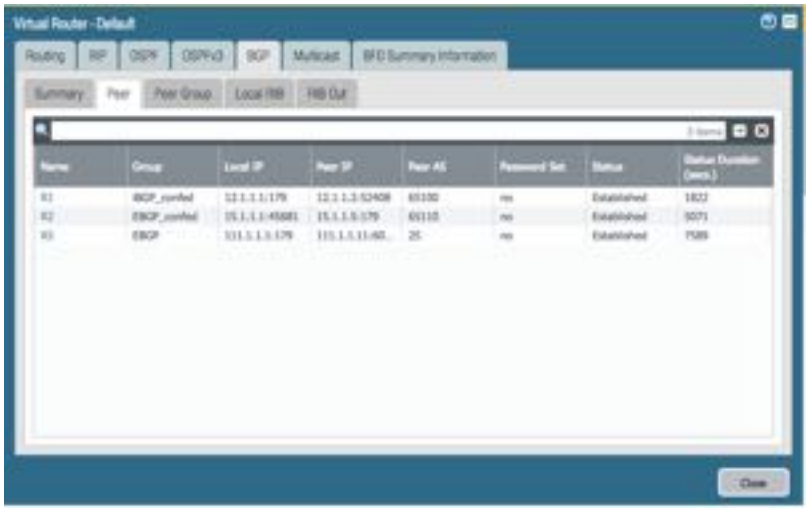




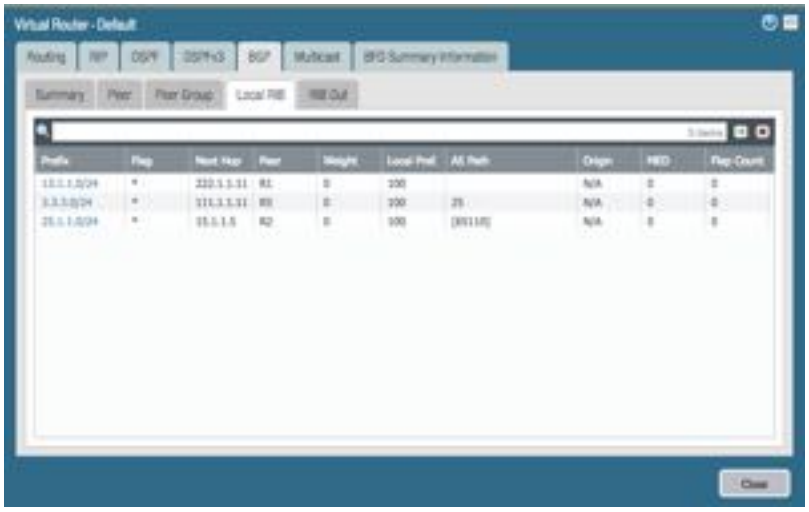
© 2018 Palo Alto Networks, Inc.



To verify that the routes from the firewall to the peers are established, on the virtual router's screen, select **More Runtime Stats** and select the **Peer** tab.

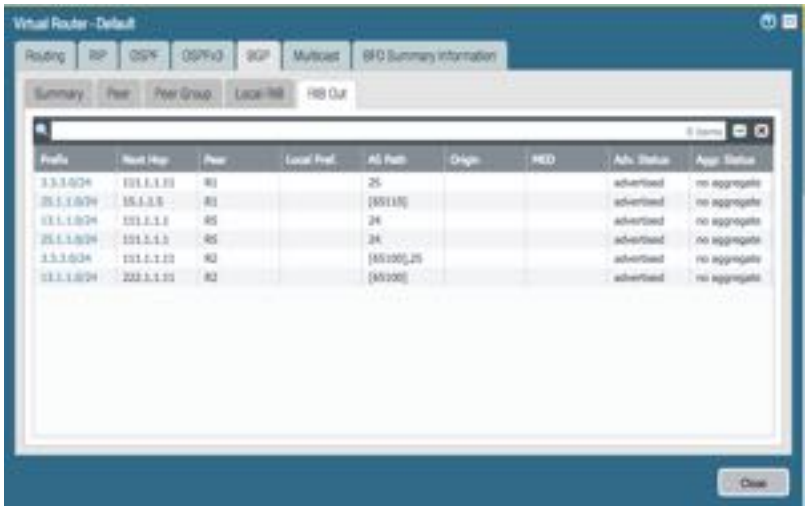


Select the **Local RIB** tab to view information about the routes stored in the Routing Information Base (RIB).



Then select the **RIB Out** tab.

© 2018 Palo Alto Networks, Inc.



© 2018 Palo Alto Networks, Inc.

IP Multicast

IP multicast is a set of protocols that network appliances use to send multicast IP datagrams to a group of interested receivers using one transmission rather than unicasting the traffic to multiple receivers, thereby saving bandwidth. IP multicast is suitable for communication from one source (or many sources) to many receivers, such as audio or video streaming, IPTV, video conferencing, and distribution of other communication, such as news and financial data.

A multicast address identifies a group of receivers that want to receive the traffic going to that address. You should not use the multicast addresses reserved for special uses, such as the range 224.0.0.0 through 224.0.0.255 or 239.0.0.0 through 239.255.255.255. Multicast traffic uses UDP, which does not resend missed packets.

Palo Alto Networks® firewalls support IP multicast and Protocol Independent Multicast (PIM) on a Layer 3 interface that you configure for a virtual router on the firewall.

For multicast routing, the Layer 3 interface type can be Ethernet, Aggregate Ethernet (AE), VLAN, loopback, or tunnel. Interface groups allow you to configure more than one firewall interface at a time with the same Internet Group Management Protocol (IGMP) and PIM parameters, and with the same group permissions (multicast groups allowed to accept traffic from any source or from only a specific source). An interface can belong to only one interface group.

The firewall supports IPv4 multicast—it does not support IPv6 multicast. The firewall also does not support PIM Dense Mode (PIM-DM), IGMP proxy, IGMP static joins, Anycast RP, GRE, or multicast configurations on a Layer 2 or virtual wire interface type. However, a virtual wire interface can pass multicast packets. Also, a Layer 2 interface can switch Layer 3 IPv4 multicast packets between different VLANs and the firewall will retag the VLAN ID using the VLAN ID of the egress interface.

You must enable multicast for a virtual router and enable PIM for an ingress and an egress interface in order for the interfaces to receive or forward multicast packets. In addition to PIM, you must also enable IGMP on egress interfaces that face receivers. You must configure a Security policy rule to allow IP multicast traffic to a predefined Layer 3 destination zone named **multicast** or to **any** destination zone.

- IGMP
- PIM
- Configure IP Multicast
- View IP Multicast Information

IGMP

Internet Group Management Protocol (IGMP) is an IPv4 protocol that a multicast receiver uses to communicate with an interface on a Palo Alto Networks® firewall and that the firewall uses to track the membership of multicast groups. When a host wants to receive multicast traffic, its implementation of IGMP sends an IGMP Membership report message and the receiving router, in turn, sends a PIM Join message to the multicast group address of the group that the host wants to join. An IGMP-enabled router on the same physical network (such as an Ethernet segment) then uses PIM to communicate with other PIM-enabled routers to determine a path from the source to interested receivers.

Enable IGMP only on interfaces that face a multicast receiver. The receivers can be only one Layer 3 hop away from the virtual router. IGMP messages are Layer 2 messages that have a TTL value of one and, therefore, cannot go outside the LAN.

When you Configure IP Multicast, specify whether an interface uses IGMP Version 1, IGMP Version 2, or IGMP Version 3. You can enforce the IP Router Alert option, RFC 2113, so that incoming IGMP packets that use IGMPv2 or IGMPv3 have the IP Router Alert option.

© 2018 Palo Alto Networks, Inc.

By default, an interface accepts IGMP Membership reports for all multicast groups. You can configure multicast group permissions to control the groups for which the virtual router accepts Membership reports from any source (Any-Source Multicast, or ASM), which is basically PIM Sparse Mode (PIM-SM). You can also specify the groups for which the virtual router accepts Membership reports from a specific source (PIM Source-Specific Multicast [PIM-SSM]). If you specify permissions for either ASM or SSM groups, the virtual router denies Membership reports from other groups. The interface must use IGMPv3 to pass PIM-SSM traffic.

You can specify the maximum number of sources and the maximum number of multicast groups that IGMP can process simultaneously for an interface.

The virtual router multicasts an IGMP Query at regular intervals to all receivers of a multicast group. A receiver responds to an IGMP Query with an IGMP Membership report that confirms the receiver still wants to receive multicast traffic for that group. The virtual router maintains a table of the multicast groups that have receivers; the virtual router forwards a multicast packet out the interface to the next hop only if there is still a receiver down that multicast distribution tree that is joined to the group. The virtual router does not track exactly which receivers are joined to a group. Only one router on a subnet responds to IGMP Queries and that is the IGMP Querier—the router with the lowest IP address.

You can configure an interface with an IGMP Query interval and the amount of time allowed for a receiver to respond to a query (the Max Query Response Time). When a virtual router receives an IGMP Leave message from a receiver to leave a group, the virtual router checks that the interface that received the Leave message is not configured with the Immediate Leave option. In the absence of the Immediate Leave option, the virtual router sends a Query to determine whether there are still receiver members for the group. The Last Member Query Interval specifies how many seconds are allowed for any remaining receivers for that group to respond and confirm that they still want multicast traffic for that group.

An interface supports the IGMP robustness variable, which you can adjust so that the firewall then tunes the Group Membership Interval, Other Querier Present Interval, Startup Query Count, and Last Member Query Count. A higher robustness variable can accommodate a subnet that is likely to drop packets.

View IP Multicast Information to see IGMP-enabled interfaces, the IGMP version, Querier address, robustness setting, limits on the number of multicast groups and sources, and whether the interface is configured for Immediate Leave. You can also see the multicast groups to which interfaces belong and other IGMP membership information.

PIM

IP multicast uses the Protocol Independent Multicast (PIM) routing protocol between routers to determine the path on the distribution tree that multicast packets take from the source to the receivers (multicast group members). A Palo Alto Networks® firewall supports PIM Sparse Mode (PIM-SM) (RFC 4601), PIM Any-Source Multicast (ASM) (sometimes referred to as PIM Sparse Mode), and PIM Source-Specific Multicast (SSM). In PIM-SM, the source does not forward multicast traffic until a receiver (user) belonging to a multicast group requests that the source send the traffic. When a host wants to receive multicast traffic, its implementation of IGMP sends an IGMP Membership report message, and the receiving router then sends a PIM Join message to the multicast group address of the group it wants to join.

- In **ASM**, the receiver uses IGMP to request traffic for a multicast group address; any source could have originated that traffic. Consequently, the receiver doesn't necessarily know the senders, and the receiver could receive multicast traffic in which it has no interest.
- In **SSM** (RFC 4607), the receiver uses IGMP to request traffic from one or more specific sources to a multicast group address. The receiver knows the IP address of the senders and receives only the multicast traffic it wants. SSM requires IGMPv3. You can override the default SSM address space, which is 232.0.0.0/8.

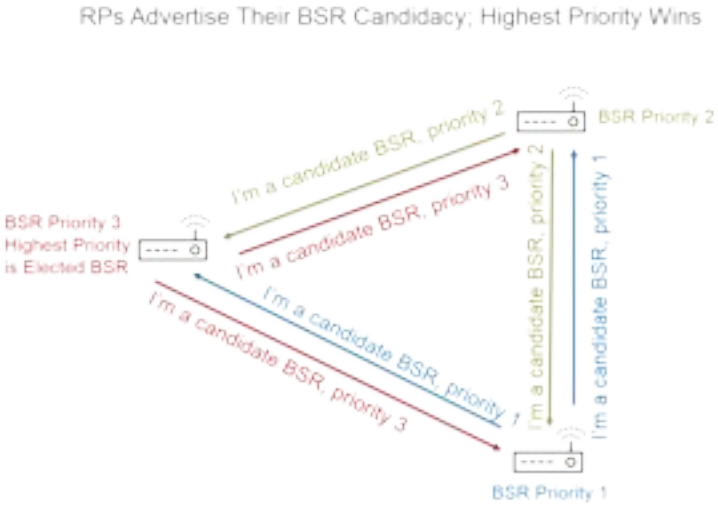
When you Configure IP Multicast on a Palo Alto Networks® firewall, you must enable PIM for an interface to forward multicast traffic, even on receiver-facing interfaces. This is unlike IGMP, which you enable only on receiver-facing interfaces.

© 2018 Palo Alto Networks, Inc.

ASM requires a *rendezvous point* (RP), which is a router located at the juncture or root of a shared distribution tree. The RP for a multicast domain serves as a single point to which all multicast group members send their Join messages. This behavior reduces the likelihood of a routing loop that would otherwise occur if group members sent their Join messages to multiple routers. (SSM doesn't need an RP because source-specific multicast uses a shortest-path tree and therefore has no need for an RP.)

In an ASM environment, there are two ways that the virtual router determines which router is the RP for a multicast group:

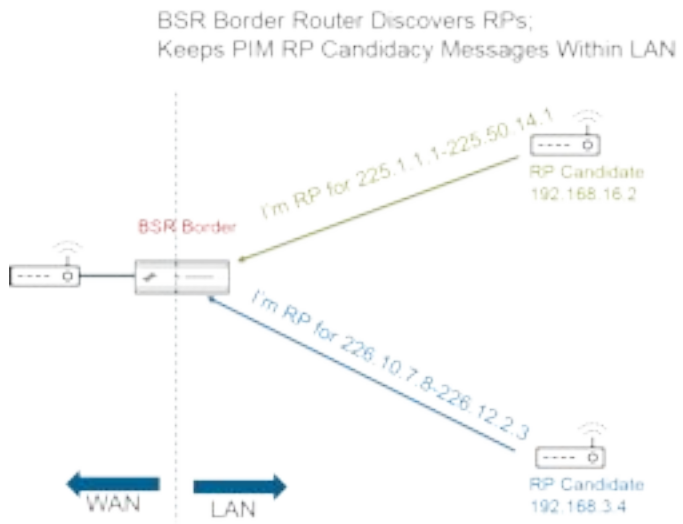
- **Static RP-to-Group Mapping**—configures the virtual router on the firewall to act as RP for multicast groups. You configure a local RP, either by configuring a static RP address or by specifying that the local RP is a candidate RP and the RP is chosen dynamically (based on lowest priority value). You can also statically configure one or more external RPs for different group address ranges not covered by the local RP, which helps you load-balance multicast traffic so that one RP is not overloaded.
- **Bootstrap Router (BSR)**—(RFC 5059)—defines the role of a BSR. First, candidates for BSR advertise their priority to each other and then the candidate with the largest priority is elected BSR, as shown in the following figure:



Next, the BSR discovers RPs when candidate RPs periodically unicast a BSR message to the BSR containing their IP address and the multicast group range for which they will act as RP. You can configure the local virtual router to be a candidate RP, in which case the virtual router announces its RP candidacy for a specific multicast group or groups. The BSR sends out RP information to the other RPs in the PIM domain.

When you configure PIM for an interface, you can select BSR Border when the interface on the firewall is at an enterprise boundary facing away from the enterprise network. The BSR Border setting prevents the firewall from sending RP candidacy BSR messages outside the LAN. In the following illustration, BSR Border is enabled for the interface facing the LAN and that interface has the highest priority. If the virtual router has both a static RP and a dynamic RP (learned from the BSR), you can specify whether the static RP should override the learned RP for a group when you configure the local, static RP.

© 2018 Palo Alto Networks, Inc.



In order for PIM Sparse Mode to notify the RP that it has traffic to send down a shared tree, the RP must be aware of the source. The host notifies the RP that it is sending traffic to a multicast group address when the *designated router* (DR) encapsulates the first packet from the host in a PIM Register message and unicasts the packet to the RP on its local network. The DR also forwards Prune messages from a receiver to the RP. The RP maintains the list of IP addresses of sources that are sending to a multicast group and the RP can forward multicast packets from sources.

Why do the routers in a PIM domain need a DR? When a router sends a PIM Join message to a switch, two routers could receive it and forward it to the same RP, causing redundant traffic and wasting bandwidth. To prevent unnecessary traffic, the PIM routers elect a DR (the router with the highest IP address), and only the DR forwards the Join message to the RP. Alternatively, you can assign a DR priority to an interface group, which takes precedence over IP address comparisons. As a reminder, the DR is forwarding (unicasting) PIM messages; it is not multicasting IP multicast packets.

You can specify the IP addresses of PIM neighbors (routers) that the interface group will allow to peer with the virtual router. By default, all PIM-enabled routers can be PIM neighbors, but the option to limit neighbors provides a step toward securing the virtual router in your PIM environment.

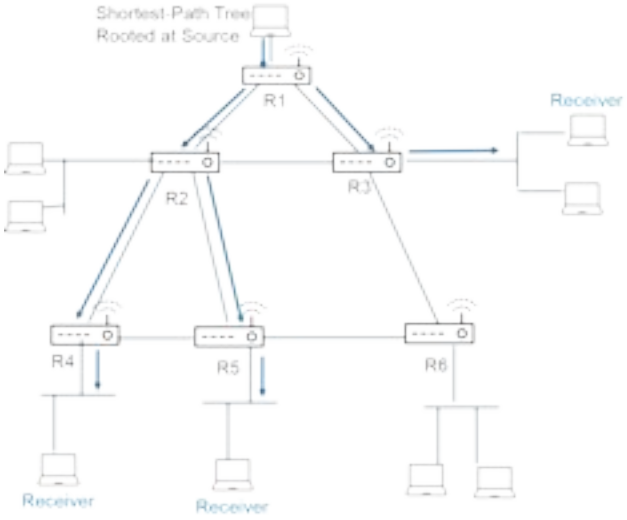
- Shortest-Path Tree (SPT) and Shared Tree
- PIM Assert Mechanism
- Reverse-Path Forwarding

Shortest-Path Tree (SPT) and Shared Tree

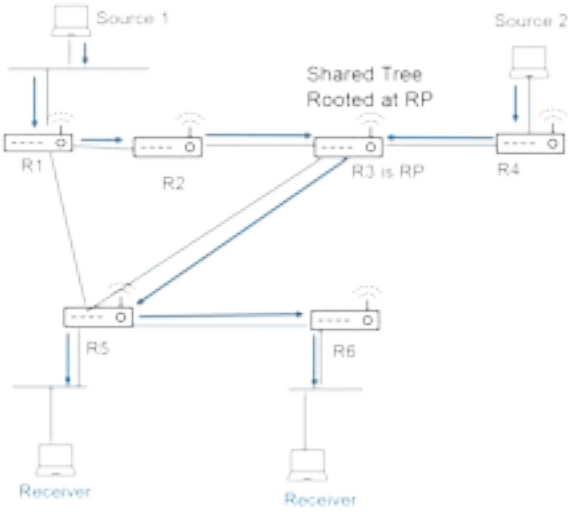
After a receiver joins a multicast group, the routers in the multiaccess network build the routing paths necessary to send data to each receiver in the group. Each IP datagram sent to a multicast group is distributed (forwarded) to all members. The routing paths constitute a type of distribution tree for a multicast packet. The goal of a multicast distribution tree is for the router to duplicate a multicast packet when the packet reaches a divergence of paths and the router must send the packet down multiple paths to reach all group members, yet the distribution tree must refrain from sending packets down a path where no interested receivers exist. The distribution tree is one of the following:

- A **source tree**—A path from a multicast source (the root of the tree) through the network to the receivers in the multicast group. The source tree is the shortest path that a multicast packet can take from source to receiver, so it is also known as the **shortest-path tree (SPT)**. The sender and receiver are annotated as a source and multicast group pair, shortened to (S, G); for example, (192.168.1.1, 225.9.2.6). The following figure illustrates three shortest-path trees from the source to three receivers.

© 2018 Palo Alto Networks, Inc.



- A **shared tree**—A path rooted at the RP, not at the multicast source. A shared tree is also known as an RP tree or RPT. Routers forward multicast packets from various sources to the RP and the RP forwards the packets down the shared tree. A shared tree is annotated as (*, G), using a wildcard as the source because all sources belonging to the multicast group share the same distribution tree from the RP. An example shared tree annotation is (*, 226.3.1.5). The following figure illustrates a shared tree from the root at the RP to the receivers.



Source-Specific Multicast (SSM) uses source tree distribution. When you Configure IP Multicast to use Any Source Multicast (ASM), you can specify which distribution tree the virtual router on your Palo Alto Networks® firewall uses to deliver multicast packets to a group by setting an SPT threshold for the group:

- By default the virtual router switches multicast routing from shared tree to SPT when it receives the first multicast packet for a group or prefix (the **SPT Threshold** is set to 0).
- You can configure the virtual router to switch to SPT when the total number of kilobits in packets arriving for the specified multicast group or prefix at any interface over any length of time reaches a configured number.
- You can configure the virtual router to never switch to SPT for the group or prefix (it continues to use shared tree).

SPT requires more memory, so choose your setting based on your multicast traffic level to the group. If the virtual router switches to SPT, then packets will arrive from the source (rather than the RP) and the virtual

© 2018 Palo Alto Networks, Inc.

router sends a Prune message to the RP. The source sends subsequent multicast packets for that group down the shortest-path tree.

PIM Assert Mechanism

To prevent routers on a multiaccess network from forwarding the same multicast traffic to the same next hop (which would cause redundant traffic and wasted bandwidth), PIM uses the Assert mechanism to elect a single PIM Forwarder for the multiaccess network.

If the virtual router receives a multicast packet from a source on an interface that the virtual router already associates as the outgoing interface for the same (S,G) pair identified in the packet, that means this is a duplicate packet. Consequently, the virtual router sends an Assert message containing its metrics to the other routers on the multiaccess network. The routers then elect a PIM Forwarder in this manner:

1. The PIM Forwarder is the router with the lowest administrative distance to the multicast source.
2. In the event of a tie for lowest administrative distance, the PIM Forwarder is the router with the best unicast routing metric to the source.
3. In the event of a tie for best metric, the PIM Forwarder is the router with the highest IP address.

Routers that are not elected as the PIM Forwarder will stop forwarding traffic to the multicast group identified in the (S,G) pair.

When you Configure IP Multicast, you can configure the interval at which the virtual router sends PIM Assert messages out an interface (the Assert interval). When you View IP Multicast Information, the **PIM Interface** tab displays the Assert interval for an interface.

Reverse-Path Forwarding

PIM uses reverse-path forwarding (RPF) to prevent multicast routing loops by leveraging the unicast routing table on the virtual router. When the virtual router receives a multicast packet, it looks up the source of the multicast packet in its unicast routing table to see if the outgoing interface associated with that source IP address is the interface on which that packet arrived. If the interfaces match, the virtual router duplicates the packet and forwards it out the interfaces toward the multicast receivers in the group. If the interfaces don't match, the virtual router drops the packet. The unicast routing table is based on the underlying static routes or the interior gateway protocol (IGP) your network uses, such as OSPF.

PIM also uses RPF to build a shortest-path tree to a source, one PIM router hop at a time. The virtual router has the address of the multicast source, so the virtual router selects as its next hop back to the source the upstream PIM neighbor that the virtual router would use to forward unicast packets to the source. The next hop router does the same thing.

After RPF succeeds and the virtual router has a route entry in its multicast routing information base (mRIB), the virtual router maintains source-based tree entries (S,G) and shared tree entries (*,G) in its multicast forwarding information base (multicast forwarding table or mFIB). Each entry includes the source IP address, multicast group, incoming interface (RPF interface) and outgoing interface list. There can be multiple outgoing interfaces for an entry because the shortest path tree can branch at the router, and the router must forward the packet out multiple interfaces to reach receivers of the group that are located down different paths. When the virtual router uses the mFIB to forward a multicast packet, it matches an (S,G) entry before it attempts to match a (*,G) entry.

If you are advertising multicast source prefixes into BGP (you configured MP-BGP with the IPv4 address family and the multicast subsequent address family), then the firewall always performs the RPF check on the BGP routes that the firewall received under the multicast subsequent address family.

View IP Multicast Information to see how to view the mFIB and mRIB entries. Keep in mind that the multicast route table (mRIB) is a separate table from the unicast route table (RIB).

© 2018 Palo Alto Networks, Inc.

Configure IP Multicast

Configure interfaces on a virtual router of a Palo Alto Networks® firewall to receive and forward IP multicast packets. You must enable IP multicast for the virtual router, configure Protocol Independent Multicast (PIM) on the ingress and egress interfaces, and configure Internet Group Management Protocol (IGMP) on receiver-facing interfaces.

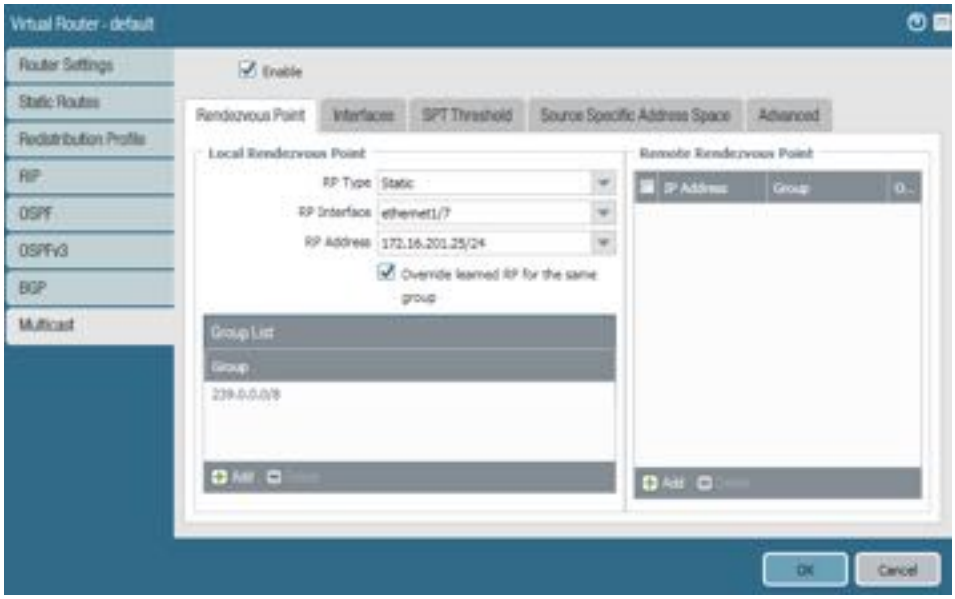
STEP 1 | Enable IP multicast for a virtual router.

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **Multicast** and **Enable** IP multicast.

STEP 2 | (ASM only) If the multicast domain in which the virtual router is located uses Any-Source Multicast (ASM), identify and configure the local and remote rendezvous points (RPs) for multicast groups.

1. Select **Rendezvous Point**.
2. Select a Local **RP Type**, which determines how the RP is chosen (the options are **Static**, **Candidate**, or **None**):

   - **Static**—Establishes a static mapping of an RP to multicast groups. Configuring a static RP requires you to explicitly configure the same RP on other PIM routers in the PIM domain.
     - Select the **RP Interface**. Valid interface types are Layer3, virtual wire, loopback, VLAN, Aggregate Ethernet (AE), and tunnel.
     - Select the **RP Address**. The IP addresses of the RP interface you selected populate the drop-down.
     - Select **Override learned RP for the same group** so that this static RP serves as RP instead of the RP elected for the groups in the Group List.
     - **Add** one or more multicast **Groups** for which the RP acts as the RP.



   - **Candidate**—Establishes a dynamic mapping of an RP to multicast groups based on priority so that each router in a PIM domain automatically elects the same RP.
     - Select the **RP Interface** of the candidate RP. Valid interface types are Layer 3, loopback, VLAN, Aggregate Ethernet (AE), and tunnel.
     - Select the **RP Address** of the candidate RP. The IP addresses for the RP interface you selected populate the drop-down.

© 2018 Palo Alto Networks, Inc.

- (Optional) Change the **Priority** for the candidate RP. The firewall compares the priority of the candidate RP to the priority of other candidate RPs to determine which one acts as RP for the specified groups; the firewall selects the candidate RP with the lowest priority value (range is 0 to 255; default is 192).
- (Optional) Change the **Advertisement Interval (sec)** (range is 1 to 26,214; default is 60).
- Enter a **Group List** of multicast groups that communicate with the RP.

- **None**—Select if this virtual router is not an RP.

3. **Add** a Remote Rendezvous Point and enter the **IP Address** of that remote (external) RP.
4. **Add** the multicast **Group Addresses** for which the specified remote RP address acts as RP.
5. Select **Override learned RP for the same group** so that the external RP you configured statically serves as RP instead of an RP that is dynamically learned (elected) for the groups in the Group Addresses list.
6. Click **OK**.

STEP 3 | Specify a group of interfaces that share a multicast configuration (IGMP, PIM, and group permissions).

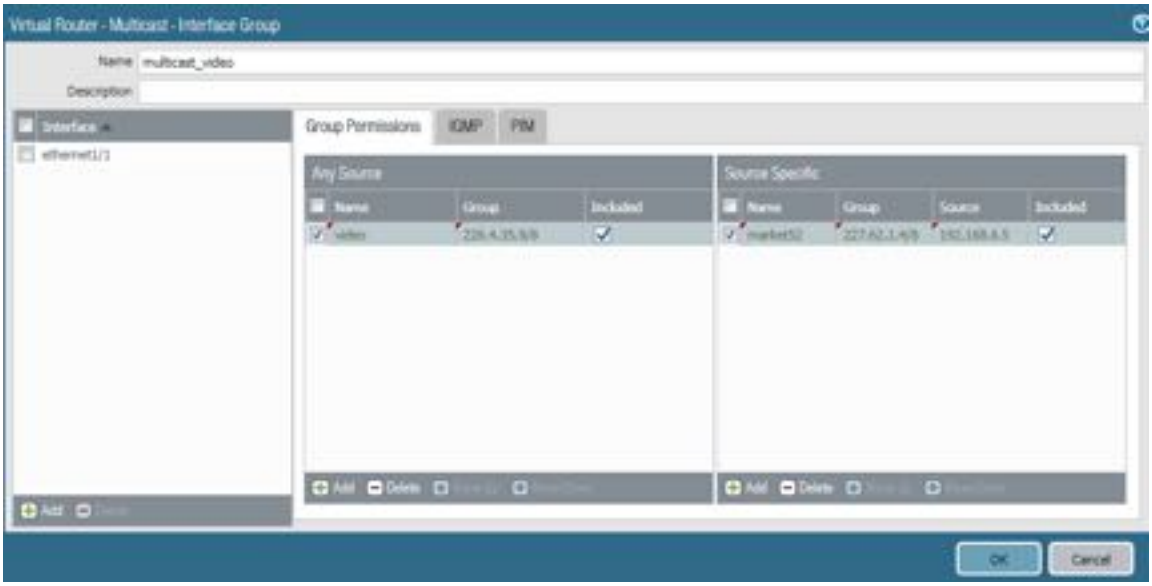
1. On the **Interfaces** tab, **Add** a **Name** for the interface group.
2. Enter a **Description**.
3. **Add** an **Interface** and select one or more Layer 3 interfaces that belong to the interface group.

STEP 4 | (Optional) Configure multicast group permissions for the interface group. By default, the interface group accepts IGMP membership reports and PIM join messages from all groups.

1. Select **Group Permissions**.
2. To configure Any-Source Multicast (ASM) groups for this interface group, in the Any Source window, **Add** a **Name** to identify a multicast group that accepts IGMP membership reports and PIM join messages from any source.
3. Enter the multicast **Group** address or group address and /prefix that can receive multicast packets from any source on these interfaces.
4. Select **Included** to include the ASM **Group** address in the interface group (default). De-select **Included** to easily exclude an ASM group from the interface group, such as during testing.
5. **Add** additional multicast **Groups** (for the interface group) that want to receive multicast packets from any source.
6. To configure Source-Specific Multicast (SSM) groups in this interface group, in the Source Specific window, **Add** a **Name** to identify a multicast group and source address pair. Don't use a name that you used for Any Source multicast. (You must use IGMPv3 to configure SSM.)
7. Enter the multicast **Group** address or group address and /prefix of the group that wants to receive multicast packets from the specified source only (and can receive the packets on these interfaces).

   A Source Specific group for which you specify permissions is a group that the virtual router must treat as source-specific. Configure Source Specific Address Space (Step 9 on page 866) that includes the source-specific groups for which you configured permission.
8. Enter the **Source** IP address from which this multicast group can receive multicast packets.
9. Select **Included** to include the SSM Group and source address pair in the interface group (default). De-select **Included** to easily exclude the pair from the interface group, such as during testing.
10. **Add** additional multicast **Groups** (for the interface group) that receive multicast packets from a specific source only.

© 2018 Palo Alto Networks, Inc.



STEP 5 | Configure IGMP for the interface group if an interface faces multicast receivers, which must use IGMP to join a group.

1. On the **IGMP** tab, **Enable** IGMP (default).
2. Specify #unique_656 parameters for interfaces in the interface group:
   - **IGMP Version**—**1**, **2**, or **3** (default).
   - **Enforce Router-Alert IP Option** (disabled by default)—Select this option if you require incoming IGMP packets that use IGMPv2 or IGMPv3 to have the IP Router Alert Option, RFC 2113.
   - **Robustness**—A variable that the firewall uses to tune the Group Membership Interval, Other Querier Present Interval, Startup Query Count, and Last Member Query Count (range is 1 to 7; default is 2). Increase the value if the subnet on which this firewall is located is prone to losing packets.
   - **Max Sources**—Maximum number of sources that IGMP can process simultaneously for an interface (range is 1 to 65,535; default is **unlimited**).
   - **Max Groups**—Maximum number of groups that IGMP can process simultaneously for an interface (range is 1 to 65,535; default is **unlimited**).
   - **Query Interval**—Number of seconds between IGMP membership Query messages that the virtual router sends to a receiver to determine whether the receiver still wants to receive the multicast packets for a group (range is 1 to 31,744; default is 125).
   - **Max Query Response Time (sec)**—Maximum number of seconds allowed for a receiver to respond to an IGMP membership Query message before the virtual router determines that the receiver no longer wants to receive multicast packets for the group (range is 0 to 3,174.4; default is 10).
   - **Last Member Query Interval (sec)**—Number of seconds allowed for a receiver to respond to a Group-Specific Query that the virtual router sends after a receiver sends a Leave Group message (range is 0.1 to 3,174.4; default is 1).
   - **Immediate Leave** (disabled by default)—When there is only one member in a multicast group and the virtual router receives an IGMP Leave message for that group, the Immediate Leave setting causes the virtual router to remove that group and outgoing interface from the multicast routing information base (mRIB) and multicast forwarding information base (mFIB) immediately, rather than waiting for the Last Member Query Interval to expire. The Immediate Leave setting saves network resources. You cannot select Immediate Leave if the interface group uses IGMPv1.

STEP 6 | Configure PIM Sparse Mode (PIM-SM) for the interface group.

1. On the **PIM** tab, **Enable** PIM (enabled by default).

© 2018 Palo Alto Networks, Inc.

2. Specify PIM parameters for the interface group:
   - **Assert Interval**—Number of seconds between #unique_657 that the virtual router sends to other PIM routers on the multiaccess network when they are electing a PIM forwarder (range is 0 to 65,534; default is 177).
   - **Hello Interval**—Number of seconds between PIM Hello messages that the virtual router sends to its PIM neighbors from each interface in the interface group (range is 0 to 18,000; default is 30).
   - **Join Prune Interval**—Number of seconds between PIM Join messages (and between PIM Prune messages) that the virtual router sends upstream toward a multicast source (range is 0 to 18,000; default is 60).
   - **DR Priority**—Designated Router (DR) priority that controls which router in a multiaccess network forwards PIM Join and Prune messages to the RP (range is 0 to 429,467,295; default is 1). The DR priority takes precedence over IP address comparisons to elect the DR.
   - **BSR Border**—Select this option if the interfaces in the interface group are on a virtual router that is the BSR located at the border of an enterprise LAN. This will prevent RP candidacy BSR messages from leaving the LAN.
3. **Add** one or more **Permitted PIM Neighbors** by specifying the **IP Address** of each router from which the virtual router accepts multicast packets.

STEP 7 | Click **OK** to save the interface group settings.

STEP 8 | (Optional) Change the Shortest-Path Tree (SPT) threshold, as described in Shortest-Path Tree (SPT) and Shared Tree.

1. Select **SPT Threshold** and **Add** a **Multicast Group/Prefix**, the multicast group or prefix for which you are specifying the distribution tree.
2. Specify the **Threshold (kb)**—The point at which routing to the specified multicast group or prefix switches from shared tree (sourced from the RP) to SPT distribution:
   - **0 (switch on first data packet)** (default)—The virtual router switches from shared tree to SPT for the group or prefix when the virtual router receives the first data packet for the group or prefix.
   - **never (do not switch to spt)**—The virtual router continues to use the shared tree to forward packets to the group or prefix.
   - Enter the total number of kilobits from multicast packets that can arrive for the multicast group or prefix at any interface and over any time period, upon which the virtual router changes to SPT distribution for that multicast group or prefix.

STEP 9 | Identify the multicast groups or groups and prefixes that accept multicast packets only from a specific source.

1. Select **Source Specific Address Space** and **Add** the **Name** for the space.
2. Enter the multicast **Group** address with prefix length to identify the address space that receives multicast packets from a specific source. If the virtual router receives a multicast packet for an SSM group but the group is not covered by a **Source Specific Address Space**, the virtual router drops the packet.
3. Select **Included** to include the source-specific address space as a multicast group address range from which the virtual router will accept multicast packets that originated from an allowed specific source. De-select **Included** to easily exclude a group address space for testing.
4. Add other source-specific address spaces to include all those groups for which you specified SSM group permission.

© 2018 Palo Alto Networks, Inc.



STEP 10 | (Optional) Change the length of time that a multicast route remains in the mRIB after the session ends between a multicast group and a source.

1. Select the **Advanced** tab.
2. Specify the **Multicast Route Age Out Time (sec)** (range is 210 to 7,200; default is 210).

STEP 11 | Click **OK** to save the multicast configuration.

STEP 12 | Create a Security policy rule to allow multicast traffic to the destination zone.

1. Create a Security Policy Rule and on the **Destination** tab, select **multicast** or **any** for the **Destination Zone**. The **multicast** zone is a predefined Layer 3 zone that matches all multicast traffic. The **Destination Address** can be a multicast group address.
2. Configure the rest of the Security policy rule.

STEP 13 | (Optional) Enable buffering of multicast packets before a route is set up.

1. Select **Device** > **Setup** > **Session** and edit Session Settings.
2. Enable **Multicast Route Setup Buffering** (disabled by default). The firewall can preserve the first packet(s) from a multicast flow if an entry for the corresponding multicast group does not yet exist in the multicast forwarding table (mFIB). The **Buffer Size** controls how many packets the firewall buffers from a flow. After the route is installed in the mFIB, the firewall automatically forwards the buffered first packet(s) to the receiver. (You need to enable multicast route setup buffering only if your content servers are directly connected to the firewall and your multicast application cannot withstand the first packet of the flow being dropped.)
3. (Optional) Change the **Buffer Size**. Buffer size is the number of packets per multicast flow that the firewall can buffer until the mFIB entry is set up (range is 1 to 2,000; default is 1,000). The firewall can buffer a maximum of 5,000 packets total (for all flows).
4. Click **OK**.

STEP 14 | **Commit** your changes.

STEP 15 | View IP Multicast Information to view mRIB and mFIB entries, IGMP interface settings, IGMP group memberships, PIM ASM and SSM modes, group mappings to RPs, DR addresses, PIM settings, PIM neighbors, and more.

STEP 16 | If you Configure a Static Route for multicast traffic, you can install the route only in the multicast routing table (not the unicast routing table) so that the route is used for multicast traffic only.

STEP 17 | If you enable IP multicast, it is not necessary to Configure BGP with MP-BGP for IPv4 Multicast unless you have a logical multicast topology separate from a logical unicast topology. You configure MP-BGP extensions with the IPv4 address family and multicast

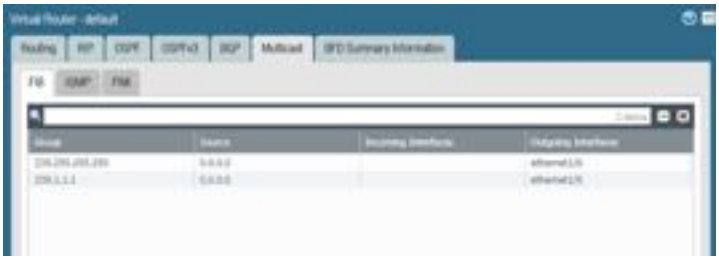
© 2018 Palo Alto Networks, Inc.

subsequent address family only when you want to advertise multicast source prefixes into BGP under multicast subsequent address family.

View IP Multicast Information

After you Configure IP Multicast routing, view multicast routes, forwarding entries, and information about your IGMP and PIM interfaces.

- Select **Network** > **Virtual Routers** and in the row for the virtual router you configured, click **More Runtime Stats**.
  1. Select **Routing** > **Route Table** and then the **Multicast** radio button to display only multicast routes (destination IP multicast group, the next hop toward that group, and outgoing interface). This information comes from the mRIB.
  2. Select **Multicast** > **FIB** to view multicast route information from the mFIB: multicast groups to which the virtual router belongs, the corresponding source, incoming interfaces, and outgoing interfaces toward the receivers.



  3. Select **Multicast** > **IGMP** > **Interface** to view IGMP-enabled interfaces, the associated IGMP version, IP address of the IGMP Querier, Querier up time and expiry time, the robustness setting, limits on numbers of multicast groups and sources, and whether the interface is configured for immediate leave.



  4. Select **Multicast** > **IGMP** > **Membership** to see IGMP-enabled interfaces and the multicast groups to which they belong, the source, and other IGMP information.



  5. Select **Multicast** > **PIM** > **Group Mapping** to view multicast groups mapped to an RP, the origin of the RP mapping, the PIM mode for the group (ASM or SSM), and whether the group is inactive. Groups in SSM mode don't use an RP, so the RP address displayed is 0.0.0.0. The default SSM group is 232.0.0.0/8.

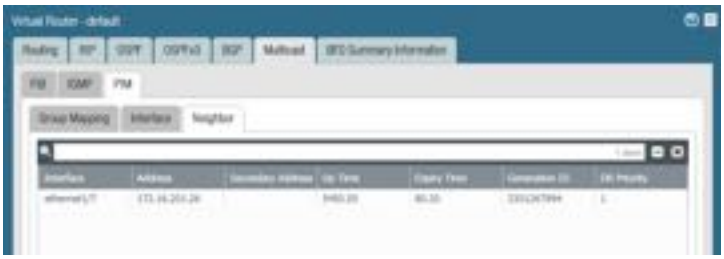
© 2018 Palo Alto Networks, Inc.



6. Select **Multicast** > **PIM** > **Interface** to view the IP address of the DR for an interface; the DR priority; the Hello, Join/Prune, and Assert intervals; and whether the interface is a bootstrap router (BSR).



7. Select **Multicast** > **PIM** > **Neighbor** to view information about routers that are PIM neighbors to the virtual router.



© 2018 Palo Alto Networks, Inc.

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 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.

This means, for example, you can make specific networks that were 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.

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 make internal OSPFv3 networks available through BGP by redistributing OSPFv3 routes into the BGP RIB.

STEP 1 | Create a redistribution profile.

1. Select **Network** > **Virtual Routers** and select a virtual router.
2. Select **Redistribution Profile** and **IPv4** or **IPv6** and **Add** a profile.
3. Enter a **Name** for the profile, which must start with an alphanumeric character and can contain zero or more underscores (_), hyphens (-), dots (.), or spaces (up to 16 characters).
4. Enter a **Priority** for the profile in the range 1 to 255. The firewall matches routes to profiles in order using the profile with the highest priority (lowest priority value) first. Higher priority rules take precedence over lower priority rules.
5. For **Redistribute**, select one of the following:
   - **Redist**—Select for redistribution the routes that match this filter.
   - **No Redist**—Select for redistribution routes that match the redistribution profiles except the routes that match this filter. This selection treats the profile as a blacklist that specifies which routes not to select for redistribution. For example, if you have multiple redistribution profiles for BGP, you can create a **No Redist** profile to exclude several prefixes, and then a general redistribution profile with a lower priority (higher priority value) after it. The two profiles combine and the higher priority profile takes precedence. You can't have only **No Redist** profiles; you would always need at least one **Redist** profile to redistribute routes.
6. On the **General Filter** tab, for Source Type, select one or more types of route to redistribute:
   - **bgp**—Redistribute BGP routes that match the profile.
   - **connect**—Redistribute connected routes that match the profile.
   - **ospf** (IPv4 only)—Redistribute OSPF routes that match the profile.
   - **rip** (IPv4 only)—Redistribute RIP routes that match the profile.
   - **ospfv3** (IPv6 only)—Redistribute OSPFv3 routes that match the profile.
   - **static**—Redistribute static routes that match the profile.
7. (Optional) For **Interface**, **Add** one or more egress interfaces of associated routes to match for redistribution. To remove an entry, click **Delete**.
8. (Optional) For **Destination**, **Add** one or more IPv4 or IPv6 destinations of routes to match for redistribution. To remove an entry, click **Delete**.

© 2018 Palo Alto Networks, Inc.

9. (Optional) For **Next Hop**, **Add** one or more next hop IPv4 or IPv6 addresses of routes to match for redistribution. To remove an entry, click **Delete**.
10. Click **OK**.

STEP 2 | (Optional—When General Filter includes ospf or ospfv3) Create an OSPF filter to further specify which OSPF or OSPFv3 routes to redistribute.

1. Select **Network** > **Virtual Routers** and select the virtual router.
2. Select **Redistribution Profile** and **IPv4** or **IPv6** and select the profile you created.
3. Select **OSPF Filter**.
4. For Path Type, select one or more of the following types of OSPF path to redistribute: **ext-1**, **ext-2**, **inter-area**, or **intra-area**.
5. To specify an **Area** from which to redistribute OSPF or OSPFv3 routes, **Add** an area in IP address format.
6. To specify a **Tag**, **Add** a tag in IP address format.
7. Click **OK**.

STEP 3 | (Optional—When General Filter includes bgp) Create a BGP filter to further specify which BGP routes to redistribute.

1. Select **Network** > **Virtual Routers** and select the virtual router.
2. Select **Redistribution Profile** and **IPv4** or **IPv6** and select the profile you created.
3. Select **BGP Filter**.
4. For **Community**, **Add** to select from the list of communities, such as well-known communities: **local-as**, **no-advertise**, **no-export**, or **nopeer**. You can also enter a 32-bit value in decimal or hexadecimal or in AS:VAL format, where AS and VAL are each in the range 0 to 65,535. Enter a maximum of 10 entries.
5. For **Extended Community**, **Add** an extended community as a 64-bit value in hexadecimal or in TYPE:AS:VAL or TYPE:IP:VAL format. TYPE is 16 bits; AS or IP is 16 bits; VAL is 32 bits. Enter a maximum of five entries.
6. Click **OK**.

STEP 4 | Select the protocol into which you are redistributing routes, and set the attributes for those routes.

This task illustrates redistributing routes into BGP.

1. Select **Network** > **Virtual Routers** and select the virtual router.
2. Select **BGP** > **Redist Rules**.
3. Select **Allow Redistribute Default Route** to allow the firewall to redistribute the default route.
4. Click **Add**.
5. Select **Address Family Type**: **IPv4** or **IPv6** to specify in which route table the redistributed routes will be put.
6. Select the **Name** of the Redistribution profile you created, which selects the routes to redistribute.
7. **Enable** the redistribution rule.
8. (Optional) Enter any of the following values, which the firewall applies to the routes being redistributed:
   - **Metric** in the range 1 to 65,535.
   - **Set Origin**—Origin of the route: **igp**, **egp**, or **incomplete**.
   - **Set MED**—MED value in the range 0 to 4,294,967,295.
   - **Set Local Preference**—Local preference value in the range 0 to 4,294,967,295.
   - **Set AS Path Limit**—Maximum number of autonomous systems in the AS_PATH in the range 1 to 255.

© 2018 Palo Alto Networks, Inc.

- **Set Community**—Select or enter a 32-bit value in decimal or hexadecimal, or enter a value in AS:VAL format, where AS and VAL are each in the range 0 to 65,525. Enter a maximum of 10 entries.
- **Set Extended Community**—Select or enter an extended community as a 64-bit value in hexadecimal or in TYPE:AS:VAL or TYPE:IP:VAL format. TYPE is 16 bits; AS or IP is 16 bits; VAL is 32 bits. Enter a maximum of five entries.

9. Click **OK**.

STEP 5 | Commit.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

DHCP

This section describes Dynamic Host Configuration Protocol (DHCP) and the tasks required to configure an interface on a Palo Alto Networks firewall to act as a DHCP server, client, or relay agent. By assigning these roles to different interfaces, the firewall can perform multiple roles.

- DHCP Overview on page 873
- Firewall as a DHCP Server and Client on page 874
- DHCP Messages on page 874
- DHCP Addressing on page 875
- DHCP Options on page 876
- Configure an Interface as a DHCP Server on page 878
- Configure an Interface as a DHCP Client on page 881
- Configure the Management Interface as a DHCP Client on page 882
- Configure an Interface as a DHCP Relay Agent on page 884
- Monitor and Troubleshoot DHCP on page 884

DHCP Overview

DHCP is a standardized protocol defined in RFC 2131, Dynamic Host Configuration Protocol. DHCP has two main purposes: to provide TCP/IP and link-layer configuration parameters and to provide network addresses to dynamically configured hosts on a TCP/IP network.

DHCP uses a client-server model of communication. This model consists of three roles that the device can fulfill: DHCP client, DHCP server, and DHCP relay agent.

- A device acting as a DHCP client (host) can request an IP address and other configuration settings from a DHCP server. Users on client devices save configuration time and effort, and need not know the network's addressing plan or other resources and options they are inheriting from the DHCP server.
- A device acting as a DHCP server can service clients. By using any of three DHCP Addressing on page 875 mechanisms, the network administrator saves configuration time and has the benefit of reusing a limited number of IP addresses when a client no longer needs network connectivity. The server can deliver IP addressing and many DHCP options to many clients.
- A device acting as a DHCP relay agent transmits DHCP messages between DHCP clients and servers.

DHCP uses User Datagram Protocol (UDP), RFC 768, as its transport protocol. DHCP messages that a client sends to a server are sent to well-known port 67 (UDP—Bootstrap Protocol and DHCP). DHCP Messages on page 874 that a server sends to a client are sent to port 68.

An interface on a Palo Alto Networks firewall can perform the role of a DHCP server, client, or relay agent. The interface of a DHCP server or relay agent must be a Layer 3 Ethernet, Aggregated Ethernet, or Layer 3 VLAN interface. You configure the firewall interfaces with the appropriate settings for any combination of roles. The behavior of each role is summarized in Firewall as a DHCP Server and Client on page 874.

The firewall supports DHCPv4 Server and DHCPv6 Relay. However, a single interface cannot support both DHCPv4 Server and DHCPv6 Relay.

The Palo Alto Networks implementations of DHCP server and DHCP client support IPv4 addresses only. Its DHCP relay implementation supports IPv4 and IPv6. DHCP client is not supported in High Availability active/active mode.

© 2018 Palo Alto Networks, Inc.

Firewall as a DHCP Server and Client

The firewall can function as a DHCP server and as a DHCP client. Dynamic Host Configuration Protocol, RFC 2131, is designed to support IPv4 and IPv6 addresses. The Palo Alto Networks implementation of DHCP server supports IPv4 addresses only.

The firewall DHCP server operates in the following manner:

- When the DHCP server receives a DHCPDISCOVER message from a client, the server replies with a DHCPOFFER message containing all of the predefined and user-defined options in the order they appear in the configuration. The client selects the options it needs and responds with a DHCPREQUEST message.
- When the server receives a DHCPREQUEST message from a client, the server replies with its DHCPACK message containing only the options specified in the request.

The firewall DHCP client operates in the following manner:

- When the DHCP client receives a DHCPOFFER from the server, the client automatically caches all of the options offered for future use, regardless of which options it had sent in its DHCPREQUEST.
- By default and to save memory consumption, the client caches only the first value of each option code if it receives multiple values for a code.
- There is no maximum length for DHCP messages unless the DHCP client specifies a maximum in option 57 in its DHCPDISCOVER or DHCPREQUEST messages.

DHCP Messages

DHCP uses eight standard message types, which are identified by an option type number in the DHCP message. For example, when a client wants to find a DHCP server, it broadcasts a DHCPDISCOVER message on its local physical subnetwork. If there is no DHCP server on its subnet and if DHCP Helper or DHCP Relay is configured properly, the message is forwarded to DHCP servers on a different physical subnet. Otherwise, the message will go no further than the subnet on which it originated. One or more DHCP servers will respond with a DHCPOFFER message that contains an available network address and other configuration parameters.

When the client needs an IP address, it sends a DHCPREQUEST to one or more servers. Of course if the client is requesting an IP address, it doesn't have one yet, so RFC 2131 requires that the broadcast message the client sends out have a source address of 0 in its IP header.

When a client requests configuration parameters from a server, it might receive responses from more than one server. Once a client has received its IP address, it is said that the client has at least an IP address and possibly other configuration parameters *bound* to it. DHCP servers manage such binding of configuration parameters to clients.

The following table lists the DHCP messages.

DHCP Message	Description
DHCPDISCOVER	Client broadcast to find available DHCP servers.
DHCPOFFER	Server response to client's DHCPDISCOVER, offering configuration parameters.
DHCPREQUEST	Client message to one or more servers to do any of the following: • Request parameters from one server and implicitly decline offers from other servers.

© 2018 Palo Alto Networks, Inc.

DHCP Message	Description
	• Confirm that a previously allocated address is correct after, for example, a system reboot. • Extend the lease of a network address.
DHCPACK	Server to client acknowledgment message containing configuration parameters, including a confirmed network address.
DHCPNAK	Server to client negative acknowledgment indicating the client's understanding of the network address is incorrect (for example, if the client has moved to a new subnet), or a client's lease has expired.
DHCPDECLINE	Client to server message indicating the network address is already being used.
DHCPRELEASE	Client to server message giving up the user of the network address and canceling the remaining time on the lease.
DHCPINFORM	Client to server message requesting only local configuration parameters; client has an externally configured network address.

DHCP Addressing

- DHCP Address Allocation Methods on page 875
- DHCP Leases on page 876

DHCP Address Allocation Methods

There are three ways that a DHCP server either assigns or sends an IP address to a client:

- **Automatic allocation**—The DHCP server assigns a permanent IP address to a client from its **IP Pools**. On the firewall, a **Lease** specified as **Unlimited** means the allocation is permanent.
- **Dynamic allocation**—The DHCP server assigns a reusable IP address from **IP Pools** of addresses to a client for a maximum period of time, known as a *lease*. This method of address allocation is useful when the customer has a limited number of IP addresses; they can be assigned to clients who need only temporary access to the network. See the DHCP Leases section.
- **Static allocation**—The network administrator chooses the IP address to assign to the client and the DHCP server sends it to the client. A static DHCP allocation is permanent; it is done by configuring a DHCP server and choosing a **Reserved Address** to correspond to the **MAC Address** of the client device. The DHCP assignment remains in place even if the client logs off, reboots, has a power outage, etc.

  Static allocation of an IP address is useful, for example, if you have a printer on a LAN and you do not want its IP address to keep changing, because it is associated with a printer name through DNS. Another example is if a client device is used for something crucial and must keep the same IP address, even if the device is turned off, unplugged, rebooted, or a power outage occurs, etc.

  Keep these points in mind when configuring a **Reserved Address**:

  - It is an address from the **IP Pools**. You may configure multiple reserved addresses.
  - If you configure no **Reserved Address**, the clients of the server will receive new DHCP assignments from the pool when their leases expire or if they reboot, etc. (unless you specified that a **Lease** is **Unlimited**).
  - If you allocate all of the addresses in the **IP Pools** as a **Reserved Address**, there are no dynamic addresses free to assign to the next DHCP client requesting an address.

© 2018 Palo Alto Networks, Inc.

- You may configure a **Reserved Address** without configuring a **MAC Address**. In this case, the DHCP server will not assign the **Reserved Address** to any device. You might reserve a few addresses from the pool and statically assign them to a fax and printer, for example, without using DHCP.

DHCP Leases

A lease is defined as the time period for which a DHCP server allocates a network address to a client. The lease might be extended (renewed) upon subsequent requests. If the client no longer needs the address, it can release the address back to the server before the lease is up. The server is then free to assign that address to a different client if it has run out of unassigned addresses.

The lease period configured for a DHCP server applies to all of the addresses that a single DHCP server (interface) dynamically assigns to its clients. That is, all of that interface's addresses assigned dynamically are of **Unlimited** duration or have the same **Timeout** value. A different DHCP server configured on the firewall may have a different lease term for its clients. A **Reserved Address** is a static address allocation and is not subject to the lease terms.

Per the DHCP standard, RFC 2131, a DHCP client does not wait for its lease to expire, because it risks getting a new address assigned to it. Instead, when a DHCP client reaches the halfway point of its lease period, it attempts to extend its lease so that it retains the same IP address. Thus, the lease duration is like a sliding window.

Typically if an IP address was assigned to a device, the device was subsequently taken off the network and its lease was not extended, the DHCP server will let that lease run out. Because the client is gone from the network and no longer needs the address, the lease duration in the server is reached and the lease is in "Expired" state.

The firewall has a hold timer that prevents the expired IP address from being reassigned immediately. This behavior temporarily reserves the address for the device in case it comes back onto the network. But if the address pool runs out of addresses, the server re-allocates this expired address before the hold timer expires. Expired addresses are cleared automatically as the systems needs more addresses or when the hold timer releases them.

In the CLI, use the **`show dhcp server lease`** operational command to view lease information about the allocated IP addresses. If you don't want to wait for expired leases to be released automatically, you can use the **`clear dhcp lease interface <interface> expired-only`** command to clear expired leases, making those addresses available in the pool again. You can use the **`clear dhcp lease interface <interface> ip <ip_address>`** command to release a particular IP address. Use the **`clear dhcp lease interface <interface> mac <mac_address>`** command to release a particular MAC address.

DHCP Options

The history of DHCP and DHCP options traces back to the Bootstrap Protocol (BOOTP). BOOTP was used by a host to configure itself dynamically during its booting procedure. A host could receive an IP address and a file from which to download a boot program from a server, along with the server's address and the address of an Internet gateway.

Included in the BOOTP packet was a vendor information field, which could contain a number of tagged fields containing various types of information, such as the subnet mask, the BOOTP file size, and many other values. RFC 1497 describes the BOOTP Vendor Information Extensions. DHCP replaces BOOTP; BOOTP is not supported on the firewall.

These extensions eventually expanded with the use of DHCP and DHCP host configuration parameters, also known as options. Similar to vendor extensions, DHCP options are tagged data items that provide information to a DHCP client. The options are sent in a variable-length field at the end of a DHCP message. For example, the DHCP Message Type is option 53, and a value of 1 indicates the DHCPDISCOVER message. DHCP options are defined in RFC 2132, DHCP Options and BOOTP Vendor Extensions.

© 2018 Palo Alto Networks, Inc.

A DHCP client can negotiate with the server, limiting the server to send only those options that the client requests.

- Predefined DHCP Options on page 877
- Multiple Values for a DHCP Option on page 877
- DHCP Options 43, 55, and 60 and Other Customized Options on page 878

Predefined DHCP Options

Palo Alto Networks firewalls support user-defined and predefined DHCP options in the DHCP server implementation. Such options are configured on the DHCP server and sent to the clients that sent a DHCPREQUEST to the server. The clients are said to *inherit* and implement the options that they are programmed to accept.

The firewall supports the following predefined options on its DHCP servers, shown in the order in which they appear on the **DHCP Server** configuration screen:

DHCP Option	DHCP Option Name
51	Lease duration
3	Gateway
1	IP Pool Subnet (mask)
6	Domain Name System (DNS on page 887) server address (primary and secondary)
44	Windows Internet Name Service (WINS) server address (primary and secondary)
41	Network Information Service (NIS) server address (primary and secondary)
42	Network Time Protocol (NTP) server address (primary and secondary)
70	Post Office Protocol Version 3 (POP3) server address
69	Simple Mail Transfer Protocol (SMTP) server address
15	DNS suffix

As mentioned, you can also configure vendor-specific and customized options, which support a wide variety of office equipment, such as IP phones and wireless infrastructure devices. Each option code supports multiple values, which can be IP address, ASCII, or hexadecimal format. With the firewall enhanced DCHP option support, branch offices do not need to purchase and manage their own DHCP servers in order to provide vendor-specific and customized options to DHCP clients.

Multiple Values for a DHCP Option

You can enter multiple option values for an **Option Code** with the same **Option Name**, but all values for a particular code and name combination must be the same type (IP address, ASCII, or hexadecimal). If one type is inherited or entered, and later a different type is entered for the same code and name combination, the second type will overwrite the first type.

© 2018 Palo Alto Networks, Inc.

You can enter an **Option Code** more than once by using a different **Option Name**. In this case, the **Option Type** for the Option Code can differ among the multiple option names. For example, if option Coastal Server (option code 6) is configured with IP address type, option Server XYZ (option code 6) with ASCII type is also allowed.

The firewall sends multiple values for an option (strung together) to a client in order from top to bottom. Therefore, when entering multiple values for an option, enter the values in the order of preference, or else move the options to achieve your preferred order in the list. The order of options in the firewall configuration determines the order that the options appear in DHCPOFFER and DHCPACK messages.

You can enter an option code that already exists as a predefined option code, and the customized option code will override the predefined DHCP option; the firewall issues a warning.

DHCP Options 43, 55, and 60 and Other Customized Options

The following table describes the option behavior for several options described in RFC 2132.

Option Code	Option Name	Option Description/Behavior
43	Vendor Specific Information	Sent from server to client. Vendor-specific information that the DHCP server has been configured to offer to the client. The information is sent to the client only if the server has a Vendor Class Identifier (VCI) in its table that matches the VCI in the client's DHCPREQUEST. An Option 43 packet can contain multiple vendor-specific pieces of information. It can also include encapsulated, vendor-specific extensions of data.
55	Parameter Request List	Sent from client to server. List of configuration parameters (option codes) that a DHCP client is requesting, possibly in order of the client's preference. The server tries to respond with options in the same order.
60	Vendor Class Identifier (VCI)	Sent from client to server. Vendor type and configuration of a DHCP client. The DHCP client sends option code 60 in a DHCPREQUEST to the DHCP server. When the server receives option 60, it sees the VCI, finds the matching VCI in its own table, and then it returns option 43 with the value (that corresponds to the VCI), thereby relaying vendor-specific information to the correct client. Both the client and server have knowledge of the VCI.

You can send custom, vendor-specific option codes that are not defined in RFC 2132. The option codes can be in the range 1-254 and of fixed or variable length.

Custom DHCP options are not validated by the DHCP Server; you must ensure that you enter correct values for the options you create.

For ASCII and hexadecimal DHCP option types, the option value can be a maximum of 255 octets.

Configure an Interface as a DHCP Server

The prerequisites for this task are:

© 2018 Palo Alto Networks, Inc.

- Configure a Layer 3 Ethernet or Layer 3 VLAN interface.
- Assign the interface to a virtual router and a zone.
- Determine a valid pool of IP addresses from your network plan that you can designate to be assigned by your DHCP server to clients.
- Collect the DHCP options, values, and Vendor Class Identifiers you plan to configure.

Perform the following task to configure an interface on the firewall to act as a DHCP server.

You can configure a combined total of 500 DHCP servers (IPv4) and DHCP relay agents (IPv4 and IPv6) on a firewall.

STEP 1 | Select an interface to be a DHCP Server.

1. Select **Network** > **DHCP** > **DHCP Server** and **Add** an **Interface** name or select one from the drop-down.
2. For **Mode**, select **enabled** or **auto** mode. Auto mode enables the server and disables it if another DHCP server is detected on the network. The **disabled** setting disables the server.
3. (Optional) Select **Ping IP when allocating new IP** if you want the server to ping the IP address before it assigns that address to its client.

   If the ping receives a response, that means a different device already has that address, so it is not available. The server assigns the next address from the pool instead. This behavior is similar to Optimistic Duplicate Address Detection (DAD)forIPv6, RFC 4429*.*

   After you set options and return to the DHCP server tab, the Probe IP column for the interface indicates if Ping IP when allocating new IP was selected.

STEP 2 | Configure the predefined DHCP Options that the server sends to its clients.

- In the Options section, select a **Lease** type:
- **Unlimited** causes the server to dynamically choose IP addresses from the **IP Pools** and assign them permanently to clients.
- **Timeout** determines how long the lease will last. Enter the number of **Days** and **Hours**, and optionally the number of **Minutes**.
- **Inheritance Source**—Leave **None** or select a source DHCP client interface or PPPoE client interface to propagate various server settings into the DHCP server. If you specify an **Inheritance Source**, select one or more options below that you want **inherited** from this source.

Specifying an inheritance source allows the firewall to quickly add DHCP options from the upstream server received by the DHCP client. It also keeps the client options updated if the source changes an option. For example, if the source replaces its NTP server (which had been identified as the **Primary NTP** server), the client will automatically inherit the new address as its **Primary NTP** server.

When inheriting DHCP option(s) that contain multiple IP addresses, the firewall uses only the first IP address contained in the option to conserve cache memory. If you require multiple IP addresses for a single option, configure the DHCP options directly on that firewall rather than configure inheritance.

- **Check inheritance source status**—If you selected an **Inheritance Source**, clicking this link opens the **Dynamic IP Interface Status** window, which displays the options that were inherited from the DHCP client.
- **Gateway**—IP address of the network gateway (an interface on the firewall) that is used to reach any device not on the same LAN as this DHCP server.

© 2018 Palo Alto Networks, Inc.

- **Subnet Mask**—Network mask used with the addresses in the **IP Pools**.

For the following fields, click the down arrow and select **None**, or **inherited**, or enter a remote server's IP address that your DHCP server will send to clients for accessing that service. If you select **inherited**, the DHCP server inherits the values from the source DHCP client specified as the **Inheritance Source**.

- **Primary DNS**, **Secondary DNS**—IP address of the preferred and alternate Domain Name System (DNS) servers.
- **Primary WINS**, **Secondary WINS**—IP address of the preferred and alternate Windows Internet Naming Service (WINS) servers.
- **Primary NIS**, **Secondary NIS**—IP address of the preferred and alternate Network Information Service (NIS) servers.
- **Primary NTP**, **Secondary NTP**—IP address of the available Network Time Protocol servers.
- **POP3 Server**—IP address of Post Office Protocol (POP3) server.
- **SMTP Server**—IP address of a Simple Mail Transfer Protocol (SMTP) server.
- **DNS Suffix**—Suffix for the client to use locally when an unqualified hostname is entered that it cannot resolve.

STEP 3 | (Optional) Configure a vendor-specific or custom DHCP option that the DHCP server sends to its clients.

1. In the Custom DHCP Options section, **Add** a descriptive **Name** to identify the DHCP option.
2. Enter the **Option Code** you want to configure the server to offer (range is 1-254). (See RFC 2132 for option codes.)
3. If the **Option Code** is **43**, the **Vendor Class Identifier** field appears. Enter a VCI, which is a string or hexadecimal value (with 0x prefix) used as a match against a value that comes from the client Request containing option 60. The server looks up the incoming VCI in its table, finds it, and returns Option 43 and the corresponding option value.
4. **Inherit from DHCP server inheritance source**—Select it only if you specified an **Inheritance Source** for the DHCP Server predefined options and you want the vendor-specific and custom options also to be **inherited** from this source.
5. **Check inheritance source status**—If you selected an **Inheritance Source**, clicking this link opens **Dynamic IP Interface Status**, which displays the options that were inherited from the DHCP client.
6. If you did not select **Inherit from DHCP server inheritance source**, select an **Option Type**: **IP Address**, **ASCII**, or **Hexadecimal**. Hexadecimal values must start with the 0x prefix.
7. Enter the **Option Value** you want the DHCP server to offer for that **Option Code**. You can enter multiple values on separate lines.
8. Click **OK**.

STEP 4 | (Optional) Add another vendor-specific or custom DHCP option.

1. Repeat the prior step to enter another custom DHCP Option.
   - You can enter multiple option values for an **Option Code** with the same **Option Name**, but all values for an **Option Code** must be the same type (**IP Address**, **ASCII**, or **Hexadecimal**). If one type is inherited or entered and a different type is entered for the same **Option Code** and the same **Option Name**, the second type will overwrite the first type.

     When entering multiple values for an option, enter the values in the order of preference, or else move the Custom DHCP Options to achieve the preferred order in the list. Select an option and click **Move Up** or **Move Down**.
   - You can enter an **Option Code** more than once by using a different **Option Name**. In this case, the **Option Type** for the Option Code can differ among the multiple option names.
2. Click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 5 | Identify the stateful pool of IP addresses from which the DHCP server chooses an address and assigns it to a DHCP client.

If you are not the network administrator for your network, ask the network administrator for a valid pool of IP addresses from the network plan that can be designated to be assigned by your DHCP server.

1. In the **IP Pools** field, **Add** the range of IP addresses from which this server assigns an address to a client. Enter an IP subnet and subnet mask (for example, 192.168.1.0/24) or a range of IP addresses (for example, 192.168.1.10-192.168.1.20).
   - An IP Pool or a **Reserved Address** is mandatory for dynamic IP address assignment.
   - An IP Pool is optional for static IP address assignment as long as the static IP addresses that you assign fall into the subnet that the firewall interface services.
2. (Optional) Repeat this step to specify another IP address pool.

STEP 6 | (Optional) Specify an IP address from the IP pools that will not be assigned dynamically. If you also specify a **MAC Address**, the **Reserved Address** is assigned to that device when the device requests an IP address through DHCP.

See the DHCP Addressing section for an explanation of allocation of a Reserved Address.

1. In the **Reserved Address** field, click **Add**.
2. Enter an IP address from the **IP Pools** (format *x.x.x.x*) that you do not want to be assigned dynamically by the DHCP server.
3. (Optional) Specify the **MAC Address** (format *xx:xx:xx:xx:xx:xx*) of the device to which you want to permanently assign the IP address you just specified.
4. (Optional) Repeat the prior two steps to reserve another address.

STEP 7 | Commit your changes.

Click **OK** and **Commit**.

Configure an Interface as a DHCP Client

Before configuring a firewall interface as a DHCP client, make sure you have configured a Layer 3 Ethernet or Layer 3 VLAN interface, and the interface is assigned to a virtual router and a zone. Perform this task if you need to use DHCP to request an IPv4 address for an interface on your firewall.

You can also Configure the Management Interface as a DHCP Client.

STEP 1 | Configure an interface as a DHCP client.

1. Select **Network** > **Interfaces**.
2. On the **Ethernet** tab or the **VLAN** tab, **Add** an interface, or select a configured interface, that you want to be a DHCP client.
3. Select the **IPv4** tab; for **Type**, select **DHCP Client**.
4. Select **Enable**.
5. (Optional) Select **Automatically create default route pointing to default gateway provided by server**. This causes the firewall to create a static route to a default gateway that will be useful when clients are trying to access many destinations that do not need to have routes maintained in a routing table on the firewall.

© 2018 Palo Alto Networks, Inc.

6. (Optional) Enter a **Default Route Metric** (priority level) for the route between the firewall and the DHCP server (range is 1-65535; there is no default metric). A route with a lower number has higher priority during route selection. For example, a route with a metric of 10 is used before a route with a metric of 100.
7. (Optional) Select **Show DHCP Client Runtime Info** to see all of the settings the client has inherited from its DHCP server.

STEP 2 | Commit your changes.

Click **OK** and **Commit**.

Now the Ethernet interface indicates **Dynamic-DHCP Client** in its **IP Address** field on the **Ethernet** tab.

STEP 3 | (Optional) See which interfaces on the firewall are configured as DHCP clients.

1. Select **Network** > **Interfaces** > **Ethernet** and look in the **IP Address** field to see which interfaces indicate DHCP Client.
2. Select **Network** > **Interfaces** > **VLAN** and look in the **IP Address** field to see which interfaces indicate DHCP Client.

Configure the Management Interface as a DHCP Client

The management interface on the firewall supports DHCP client for IPv4, which allows the management interface to receive its IPv4 address from a DHCP server. The management interface also supports DHCP Option 12 and Option 61, which allow the firewall to send its hostname and client identifier, respectively, to DHCP servers.

By default, VM-Series firewalls deployed in AWS and Azure™ use the management interface as a DHCP client to obtain its IP address, rather than a static IP address, because cloud deployments require the automation this feature provides. DHCP on the management interface is turned off by default for the VM-Series firewall except for the VM-Series firewall in AWS and Azure. The management interfaces on WildFire and Panorama models do not support this DHCP functionality.



- *For hardware-based firewall models (not VM-Series), configure the management interface with a static IP address when possible.*
- *If the firewall acquires a management interface address through DHCP, assign a MAC address reservation on the DHCP server that serves that firewall. The reservation ensures that the firewall retains its management IP address after a restart. If the DHCP server is a Palo Alto Networks firewall, see Step 6 of* Configurean Interface as a DHCP Server *for reserving an address.*

If you configure the management interface as a DHCP client, the following restrictions apply:

- You cannot use the management interface in an HA configuration for control link (HA1 or HA1 backup), data link (HA2 or HA2 backup), or packet forwarding (HA3) communication.
- You cannot select **MGT** as the Source Interface when you customize service routes (**Device** > **Setup** > **Services** > **Service Route Configuration** > **Customize**). However, you can select **Use default** to route the packets via the management interface.
- You cannot use the dynamic IP address of the management interface to connect to a Hardware Security Module (HSM). The IP address on the HSM client firewall must be a static IP address because HSM authenticates the firewall using the IP address, and operations on HSM would stop working if the IP address were to change during runtime.

A prerequisite for this task is that the management interface must be able to reach a DHCP server.

© 2018 Palo Alto Networks, Inc.

STEP 1 | Configure the Management interface as a DHCP client so that it can receive its IP address (IPv4), netmask (IPv4), and default gateway from a DHCP server.

Optionally, you can also send the hostname and client identifier of the management interface to the DHCP server if the orchestration system you use accepts this information.

1. Select **Device** > **Setup** > **Management** and edit Management Interface Settings.
2. For **IP Type**, select **DHCP Client**.
3. (Optional) Select one or both options for the firewall to send to the DHCP server in DHCP Discover or Request messages:
   - **Send Hostname**—Sends the **Hostname** (as defined in **Device** > **Setup** > **Management**) as part of DHCP Option 12.
   - **Send Client ID**—Sends the client identifier as part of DHCP Option 61. A client identifier uniquely identifies a DHCP client, and the DHCP Server uses it to index its configuration parameter database.
4. Click **OK**.

STEP 2 | (Optional) Configure the firewall to accept the host name and domain from the DHCP server.

1. Select **Device** > **Setup** > **Management** and edit General Settings.
2. Select one or both options:
   - **Accept DHCP server provided Hostname**—Allows the firewall to accept the hostname from the DHCP server (if valid). When enabled, the hostname from the DHCP server overwrites any existing **Hostname** specified in **Device** > **Setup** > **Management**. Don't select this option if you want to manually configure a hostname.
   - **Accept DHCP server provided Domain**—Allows the firewall to accept the domain from the DHCP Server. The domain (DNS suffix) from the DHCP Server overwrites any existing **Domain** specified in **Device** > **Setup** > **Management**. Don't select this option if you want to manually configure a domain.
3. Click **OK**.

STEP 3 | Commit your changes.

Click **Commit**.

STEP 4 | View DHCP client information.

1. Select **Device** > **Setup** > **Management** and Management Interface Settings.
2. Click **Show DHCP Client Runtime Info**.

STEP 5 | (Optional) Renew the DHCP lease with the DHCP server, regardless of the lease term.

This option is convenient if you are testing or troubleshooting network issues.

1. Select **Device** > **Setup** > **Management** and edit Management Interface Settings.
2. Click **Show DHCP Client Runtime Info**.
3. Click **Renew**.

STEP 6 | (Optional) Release the following DHCP options that came from the DHCP server:

- IP Address
- Netmask
- Default Gateway
- DNS Server (primary and secondary)
- NTP Server (primary and secondary)
- Domain (DNS Suffix)

© 2018 Palo Alto Networks, Inc.

A release frees the IP address, which drops your network connection and renders the firewall unmanageable if no other interface is configured for management access.

Use the CLI operational command **`request dhcp client management-interface release`**.

Configure an Interface as a DHCP Relay Agent

To enable a firewall interface to transmit DHCP messages between clients and servers, you must configure the firewall as a DHCP relay agent. The interface can forward messages to a maximum of eight external IPv4 DHCP servers and eight external IPv6 DHCP servers. A client DHCPDISCOVER message is sent to all configured servers, and the DHCPOFFER message of the first server that responds is relayed back to the requesting client.

You can configure a combined total of 500 DHCP servers (IPv4) and DHCP relay agents (IPv4 and IPv6) on a firewall.

Before configuring a DHCP relay agent, make sure you have configured a Layer 3 Ethernet or Layer 3 VLAN interface, and the interface is assigned to a virtual router and a zone.

STEP 1 | Select DHCP Relay.

Select **Network** > **DHCP** > **DHCP Relay**.

STEP 2 | Specify the IP address of each DHCP server with which the DHCP relay agent will communicate.

1. In the **Interface** field, select from the drop-down the interface you want to be the DHCP relay agent.
2. Select either **IPv4** or **IPv6**, indicating the type of DHCP server address you will specify.
3. If you checked **IPv4**, in the **DHCP Server IP Address** field, **Add** the address of the DHCP server to and from which you will relay DHCP messages.
4. If you checked **IPv6**, in the **DHCP Server IPv6 Address** field, **Add** the address of the DHCP server to and from which you will relay DHCP messages. If you specify a *multicast* address, also specify an outgoing **Interface**.
5. (Optional) Repeat the prior three steps to enter a maximum of eight DHCP server addresses per IP address family.

STEP 3 | Commit the configuration.

Click **OK** and **Commit**.

Monitor and Troubleshoot DHCP

You can view the status of dynamic address leases that your DHCP server has assigned or that your DHCP client has been assigned by issuing commands from the CLI. You can also clear leases before they time out and are released automatically.

- View DHCP Server Information on page 885
- Clear DHCP Leases on page 885
- View DHCP Client Information on page 886
- Gather Debug Output about DHCP on page 886

© 2018 Palo Alto Networks, Inc.

View DHCP Server Information

Perform this task to view DHCP pool statistics, IP addresses the DHCP server has assigned, the corresponding MAC address, state and duration of the lease, and time the lease began. If the address was configured as a **Reserved Address**, the `state` column indicates `reserved` and there is no `duration` or `lease_time`. If the lease was configured as **Unlimited**, the `duration` column displays a value of `0`.

- View DHCP pool statistics, IP address the DHCP server assigned, MAC address, state and duration of lease, and lease start time.

```
admin@PA-200> show dhcp server lease interface all
```

```
interface: "ethernet1/2"
Allocated IPs: 1, Total number of IPs in pool: 5. 20.0000% used
ip              mac               state      duration    lease_time
192.168.3.11    f0:2f:af:42:70:cf  committed  0           Wed Jul  2
 08:10:56 2014
admin@PA-200>
```

- View the options that a DHCP server has assigned to clients.

```
admin@PA-200> show dhcp server settings all
```

```
Interface    GW           DNS1        DNS2       DNS-Suffix Inherit source
------------------------------------------------------------------------------
ethernet1/2  192.168.3.1  10.43.2.10  10.44.2.10             ethernet1/3
admin@PA-200>
```

Clear DHCP Leases

You have several options for clearing DHCP leases.

- Release expired DHCP Leases of an interface (server), such as ethernet1/2, before the hold timer releases them automatically. Those addresses will be available in the IP pool again.

```
admin@PA-200> clear dhcp lease interface ethernet1/2 expired-only
```

- Release the lease of a particular IP address, for example, 192.168.3.1.

```
admin@PA-200> clear dhcp lease interface ethernet1/2 ip 192.168.3.1
```

- Release the lease of a particular MAC address, for example, f0:2c:ae:29:71:34.

```
admin@PA-200> clear dhcp lease interface ethernet1/2 mac f0:2c:ae:29:71:34
```


© 2018 Palo Alto Networks, Inc.

View DHCP Client Information

To view the status of IP address leases sent to the firewall when it is acting as a DHCP client, use either of these CLI commands.

- **admin@PA-200>show dhcp client state *<interface_name>***
- **admin@PA-200> show dhcp client state all**

```
Interface       State        IP              Gateway          Leased-until
--------------------------------------------------------------------------
ethernet1/1     Bound        10.43.14.80     10.43.14.1       70315
admin@PA-200>
```

Gather Debug Output about DHCP

To gather debug output about DHCP, use one of the following commands:

- **admin@PA-200> debug dhcpd**
- **admin@PA-200> debug management-server dhcpd**

© 2018 Palo Alto Networks, Inc.

DNS

Domain Name System (DNS) is a protocol that translates (resolves) a user-friendly domain name, such as www.paloaltonetworks.com, to an IP address so that users can access computers, websites, services, or other resources on the internet or private networks.

- DNS Overview on page 887
- DNS Proxy Object on page 888
- DNS Server Profile on page 888
- Multi-Tenant DNS Deployments on page 889
- Configure a DNS Proxy Object on page 890
- Configure a DNS Server Profile on page 891
- Use Case 1: Firewall Requires DNS Resolution for Management Purposes on page 892
- Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System on page 893
- Use Case 3: Firewall Acts as DNS Proxy Between Client and Server on page 895
- DNS Proxy Rule and FQDN Matching on page 896

DNS Overview

DNS performs a crucial role in enabling user access to network resources so that users need not remember IP addresses and individual computers need not store a huge volume of domain names mapped to IP addresses. DNS employs a client/server model; a DNS server resolves a query for a DNS client by looking up the domain in its cache and if necessary sending queries to other servers until it can respond to the client with the corresponding IP address.

The DNS structure of domain names is hierarchical; the top-level domain (TLD) in a domain name can be a generic TLD (gTLD): com, edu, gov, int, mil, net, or org (gov and mil are for the United States only) or a country code (ccTLD), such as au (Australia) or us (United States). ccTLDs are generally reserved for countries and dependent territories.

A fully qualified domain name (FQDN) includes at a minimum a host name, a second-level domain, and a TLD to completely specify the location of the host in the DNS structure. For example, www.paloaltonetworks.com is an FQDN.

Wherever a Palo Alto Networks firewall uses an FQDN in the user interface or CLI, the firewall must resolve that FQDN using DNS. Depending on where the FQDN query originates, the firewall determines which DNS settings to use to resolve the query. The following firewall tasks are related to DNS:

- Configure your firewall with at least one DNS server so it can resolve hostnames. Configure primary and secondary DNS servers or a DNS Proxy object that specifies such servers, as shown in Use Case 1: Firewall Requires DNS Resolution for Management Purposes.
- Customize how the firewall handles DNS resolution initiated by Security policy rules, reporting, and management services (such as email, Kerberos, SNMP, syslog, and more) for each virtual system, as shown in Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System.
- Configure the firewall to act as a DNS server for a client, as shown in Use Case 3: Firewall Acts as DNS Proxy Between Client and Server.
- Configure an Anti-Spyware profile to Use DNS Queries to Identify Infected Hosts on the Network.
- Enable Passive DNS Monitoring, which allows the firewall to automatically share domain-to-IP address mappings based on your network traffic with Palo Alto Networks. The Palo Alto Networks threat research team uses this information to gain insight into malware propagation and evasion techniques that abuse the DNS system.
- Enable Evasion Signatures and then enable evasion signatures for threat prevention.

© 2018 Palo Alto Networks, Inc.

- Configure an Interface as a DHCP Server. This enables the firewall to act as a DHCP Server and sends DNS information to its DHCP clients so the provisioned DHCP clients can reach their respective DNS servers.

DNS Proxy Object

When configured as a DNS proxy, the firewall is an intermediary between DNS clients and servers; it acts as a DNS server itself by resolving queries from its DNS proxy cache. If it doesn't find the domain name in its DNS proxy cache, the firewall searches for a match to the domain name among the entries in the specific DNS proxy object (on the interface on which the DNS query arrived). The firewall forwards the query to the appropriate DNS server based on the match results. If no match is found, the firewall uses default DNS servers.

A DNS proxy object is where you configure the settings that determine how the firewall functions as a DNS proxy. You can assign a DNS proxy object to a single virtual system or it can be shared among all virtual systems.

- If the DNS proxy object is for a virtual system, you can specify a DNS Server Profile on page 888, which specifies the primary and secondary DNS server addresses, along with other information. The DNS server profile simplifies configuration.
- If the DNS proxy object is shared, you must specify at least the primary address of a DNS server.



When configuring multiple tenants (ISP subscribers) with DNS services, each tenant should have its own DNS proxy defined, which keeps the tenant's DNS service separate from other tenants' services.

In the proxy object, you specify the interfaces for which the firewall is acting as DNS proxy. The DNS proxy for the interface does not use the service route; responses to the DNS requests are always sent to the interface assigned to the virtual router where the DNS request arrived.

When you Configure a DNS Proxy Object on page 890, you can supply the DNS proxy with static FQDN-to-address mappings. You can also create DNS proxy rules that control to which DNS server the domain name queries (that match the proxy rules) are directed. You can configure a maximum of 256 DNS proxy objects on a firewall.

When the firewall receives an FQDN query (and the domain name is not in the DNS proxy cache), the firewall compares the domain name from the FQDN query to the domain names in DNS Proxy rules of the DNS Proxy object. If you specify multiple domain names in a single DNS Proxy rule, a query that matches any one of the domain names in the rule means the query matches the rule. DNS Proxy Rule and FQDN Matching on page 896 describes how the firewall determines whether an FQDN matches a domain name in a DNS proxy rule. A DNS query that matches a rule is sent to the primary DNS server configured for the proxy object to be resolved.

DNS Server Profile

To simplify configuration for a virtual system, a DNS server profile allows you to specify the virtual system that is being configured, an inheritance source or the primary and secondary IP addresses for DNS servers, and a source interface and source address (service route) that will be used in packets sent to the DNS server. The source interface determines the virtual router, which has a route table. The destination IP address is looked up in the route table of the virtual router where the source interface is assigned. It's possible that the result of the destination IP egress interface differs from the source interface. The packet would egress out of the destination IP egress interface determined by the route table lookup, but the source IP address would be the address configured. The source address is used as the destination address in the reply from the DNS server.

The virtual system report and virtual system server profile send their queries to the DNS server specified for the virtual system, if there is one. (The DNS server used is defined in **Device** > **Virtual Systems** >

© 2018 Palo Alto Networks, Inc.

General > **DNS Proxy**.) If there is no DNS server specified for the virtual system, the DNS server specified for the firewall is queried.

You Configure a DNS Server Profile on page 891 for a virtual system only; it is not for a global Shared location.

Multi-Tenant DNS Deployments

The firewall determines how to handle DNS requests based on where the request originated. An environment where an ISP has multiple tenants on a firewall is known as multi-tenancy. There are three use cases for multi-tenant DNS deployments:

- **Global Management DNS Resolution**—The firewall needs DNS resolution for its own purposes, for example, the request comes from the management plane to resolve an FQDN for a management event such as a software update service. The firewall uses the service route to get to a DNS server because DNS request isn't coming in on a specific virtual router.
- **Policy and Report FQDN Resolution for a Virtual System**—For DNS queries from a security policy, a report, or a service, you can specify a set of DNS servers specific to the virtual system (tenant) or you can default to the global DNS servers. If your use case requires a different set of DNS servers per virtual system, you must configure a DNS Proxy Object. The resolution is specific to the virtual system to which the DNS proxy is assigned. If you don't have specific DNS servers applicable to this virtual system, the firewall uses the global DNS settings.
- **Dataplane DNS Resolution for a Virtual System**—This method is also known as a Network Request for DNS Resolution. The tenant's virtual system can be configured so that specified domain names are resolved on the tenant's DNS server in its network. This method supports *split DNS*, meaning that the tenant can also use its own ISP DNS servers for the remaining DNS queries not resolved on its own server. DNS Proxy Object rules control the split DNS; the tenant's domain redirects DNS requests to its DNS servers, which are configured in a DNS server profile. The DNS server profile has primary and secondary DNS servers designated, and also DNS service routes for IPv4 and IPv6, which override the default DNS settings.

The following table summarizes the DNS resolution types. The binding location determines which DNS proxy object is used for the resolution. For illustration purposes, the use cases show how a service provider might configure DNS settings to provide DNS services for resolving DNS queries required on the firewall and for tenant (subscriber) virtual systems.

Resolution Type	Location: Shared	Location: Specific Vsys
Firewall DNS resolution—performed by management plane	Binding: Global Illustrated in Use Case 1	N/A
Security profile, reporting, and server profile resolution—performed by management plane	Binding: Global Same behavior as Use Case 1	Binding: Specific vsys Illustrated in Use Case 2
DNS proxy resolution for DNS client hosts connected to interface on firewall, going through the firewall to a DNS Server—performed by dataplane	Binding: Interface Service Route: Interface and IP address on which the DNS Request was received. Illustrated in Use Case 3	

- Use Case 1: Firewall Requires DNS Resolution for Management Purposes
- Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System

© 2018 Palo Alto Networks, Inc.

- Use Case 3: Firewall Acts as DNS Proxy Between Client and Server

Configure a DNS Proxy Object

If your firewall is to act as a DNS proxy, perform this task to configure a DNS Proxy Object. The proxy object can either be shared among all virtual systems or applied to a specific virtual system.



When the firewall is enabled to act as a DNS proxy, evasion signatures that detected crafted HTTP or TLS requests can alert to instances where a client connects to a domain other than the domains specified in the original DNS query. As a best practice, Enable Evasion Signatures *after configuring DNS proxy.*

STEP 1 | Configure the basic settings for a DNS Proxy object.

1. Select **Network** > **DNS Proxy** and **Add** a new object.
2. Verify that **Enable** is selected.
3. Enter a **Name** for the object.
4. For **Location**, select the virtual system to which the object applies. If you select **Shared**, you must specify at least a **Primary** DNS server address, and optionally a **Secondary** address.
5. If you selected a virtual system, for **Server Profile**, select a DNS Server profile or else click **DNS Server Profile** to configure a new profile. See Configure a DNS Server Profile.
6. For Inheritance Source, select a source from which to inherit default DNS server settings. The default is **None**.
7. For **Interface**, click **Add** and specify the interfaces to which the DNS Proxy object applies.
   - If you use the DNS Proxy object for performing DNS lookups, an interface is required. The firewall will listen for DNS requests on this interface, and then proxy them.
   - If you use the DNS Proxy object for a service route, the interface is optional.

STEP 2 | (Optional) Specify DNS Proxy rules.

1. On the **DNS Proxy Rules** tab, **Add** a **Name** for the rule.
2. **Turn on caching of domains resolved by this mapping** if you want the firewall to cache the resolved domains.
3. For **Domain Name**, **Add** one or more domains, one entry per row, to which the firewall compares FQDN queries. If a query matches one of the domains in the rule, the query is sent to one of the following servers to be resolved (depending on what you configured in the prior step):
   - The **Primary** or **Secondary** DNS Server directly specified for this proxy object.
   - The **Primary** or **Secondary** DNS Server specified in the DNS Server profile for this proxy object.

   DNS Proxy Rule and FQDN Matching describes how the firewall matches domain names in an FQDN to a DNS proxy rule. If no match is found, default DNS servers resolve the query.
4. Do one of the following, depending on what you set the **Location** to:
   - If you chose a virtual system, select a **DNS Server profile**.
   - If you chose **Shared**, enter a **Primary** and optionally a **Secondary** address.
5. Click **OK**.

STEP 3 | (Optional) Supply the DNS Proxy with static FQDN-to-address entries. Static DNS entries allow the firewall to resolve the FQDN to an IP address without sending a query to the DNS server.

1. On the **Static Entries** tab, **Add** a **Name**.
2. Enter the Fully Qualified Domain Name (**FQDN**).
3. For **Address**, **Add** the IP address to which the FQDN should be mapped.

© 2018 Palo Alto Networks, Inc.

You can provide additional IP addresses for an entry. The firewall will provide all of the IP addresses in its DNS response and the client chooses which address to use.

4. Click **OK**.

STEP 4 | (Optional) Enable caching and configure other advanced settings for the DNS Proxy.

1. On the **Advanced** tab, select **TCP Queries** to enable DNS queries using TCP.
   - **Max Pending Requests**—Enter the maximum number of concurrent, pending TCP DNS requests that the firewall will support (range is 64-256; default is 64).
2. For **UDP Queries Retries**, enter:
   - **Interval (sec)**—The length of time (in seconds) after which another request is sent if no response has been received (range is 1-30; default is 2).
   - **Attempts**—The maximum number of UDP query attempts (excluding the first attempt) after which the next DNS server is queried (range is 1-30; default is 5.)
3. Select **Cache** to enable the firewall to cache FQDN-to-address mappings that it learns.
   - Select **Enable TTL** to limit the length of time the firewall caches DNS resolution entries for the proxy object. Disabled by default.
     - Enter **Time to Live (sec)**, the number of seconds after which all cached entries for the proxy object are removed. After the entries are removed, new DNS requests must be resolved and cached again. Range is 60-86,400. There is no default TTL; entries remain until the firewall runs out of cache memory.
   - **Cache EDNS Responses**—Select this if you want the firewall to cache partial DNS responses that are greater than 512 bytes. If a subsequent FQDN for a cached entry arrives, the firewall sends the partial DNS response. If you want full DNS responses (greater than 512 bytes), don't select this option.

STEP 5 | Commit your changes.

Click **OK** and **Commit**.

Configure a DNS Server Profile

Configure a DNS Server Profile, which simplifies configuration of a virtual system. The **Primary DNS** or **Secondary DNS** address is used to create the DNS request that the virtual system sends to the DNS server.

STEP 1 | Name the DNS server profile, select the virtual system to which it applies, and specify the primary and secondary DNS server addresses.

1. Select **Device** > **Server Profiles** > **DNS** and **Add** a **Name** for the DNS server profile.
2. For **Location**, select the virtual system to which the profile applies.
3. For **Inheritance Source**, from the drop-down, select **None** if the DNS server addresses are not inherited. Otherwise, specify the DNS server from which the profile should inherit settings. If you choose a DNS server, click **Check inheritance source status** to see that information.
4. Specify the IP address of the **Primary DNS** server, or leave as **inherited** if you chose an **Inheritance Source**.

   Keep in mind that if you specify an FQDN instead of an IP address, the DNS for that FQDN is resolved in Device > Virtual Systems > DNS Proxy.

5. Specify the IP address of the **Secondary DNS** server, or leave as **inherited** if you chose an **Inheritance Source**.

STEP 2 | Configure the service route that the firewall automatically uses, based on whether the target DNS Server has an IP address family type of IPv4 or IPv6.

© 2018 Palo Alto Networks, Inc.

1. Click **Service Route IPv4** to enable the subsequent interface and IPv4 address to be used as the service route, if the target DNS address is an IPv4 address.
2. Specify the **Source Interface** to select the DNS server's source IP address that the service route will use. The firewall determines which virtual router is assigned that interface, and then does a route lookup in the virtual router routing table to reach the destination network (based on the **Primary DNS** address).
3. Specify the IPv4 **Source Address** from which packets going to the DNS server are sourced.
4. Click **Service Route IPv6** to enable the subsequent interface and IPv6 address to be used as the service route, if the target DNS address is an IPv6 address.
5. Specify the **Source Interface** to select the DNS server's source IP address that the service route will use. The firewall determines which virtual router is assigned that interface, and then does a route lookup in the virtual router routing table to reach the destination network (based on the **Primary DNS** address).
6. Specify the IPv6 **Source Address** from which packets going to the DNS server are sourced.
7. Click **OK**.

STEP 3 | Commit the configuration.

Click **OK** and **Commit**.

Use Case 1: Firewall Requires DNS Resolution for Management Purposes

In this use case, the firewall is the client requesting DNS resolutions of FQDNs for management events such as software update services, dynamic software updates, or WildFire. The shared, global DNS services perform the DNS resolution for the management plane functions.



STEP 1 | Configure the primary and secondary DNS servers you want the firewall to use for its management DNS resolutions.

You must manually configure at least one DNS server on the firewall or it won't be able to resolve hostnames; it won't use DNS server settings from another source, such as an ISP.

1. Select **Device** > **Setup** > **Services** > **Global** and Edit. (For firewalls that do not support multiple virtual systems, there is no **Global** tab; simply edit the Services.)
2. On the **Services** tab, for **DNS**, click **Servers** and enter the **Primary DNS Server** address and **Secondary DNS Server** address.
3. Click **OK** and **Commit**.

STEP 2 | Alternatively, you can configure a DNS Proxy Object if you want to configure advanced DNS functions such as split DNS, DNS proxy overrides, DNS proxy rules, static entries, or DNS inheritance.

1. Select **Device** > **Setup** > **Services** > **Global** and Edit.
2. On the **Services** tab, for **DNS**, select **DNS Proxy Object**.

© 2018 Palo Alto Networks, Inc.

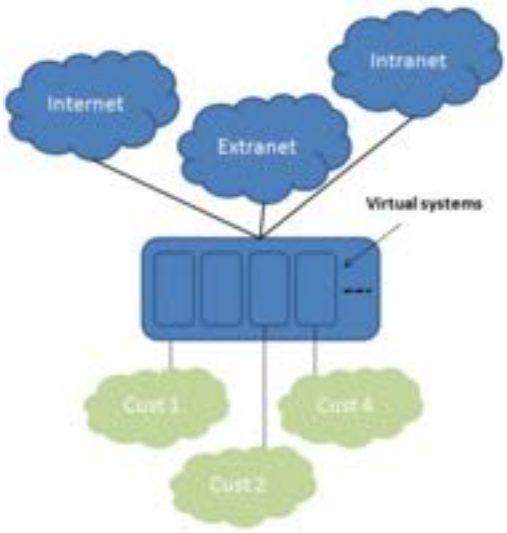
3. From the **DNS Proxy** drop-down, select the DNS proxy that you want to use to configure global DNS services, or click **DNS Proxy** to configure a new DNS proxy object as follows:
    1. Click **Enable** and enter a **Name** for the DNS proxy object.
    2. For **Location**, select **Shared** for global, firewall-wide DNS proxy services.

    Shared DNS proxy objects don't use DNS server profiles because they don't require a specific service route belonging to a tenant virtual system.

    3. Enter the **Primary** DNS server IP address. Optionally enter a **Secondary** DNS server IP address.
4. Click **OK** and **Commit**.

Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System

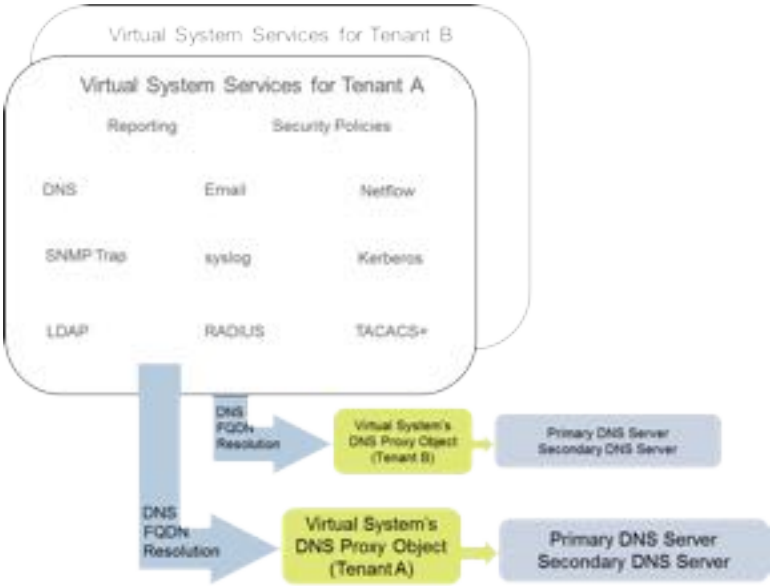
In this use case, multiple tenants (ISP subscribers) are defined on the firewall and each tenant is allocated a separate virtual system (vsys) and virtual router in order to segment its services and administrative domains. The following figure illustrates several virtual systems within a firewall.



Each tenant has its own server profiles for Security policy rules, reporting, and management services (such as email, Kerberos, SNMP, syslog, and more) defined in its own networks.

For the DNS resolutions initiated by these services, each virtual system is configured with its own DNS Proxy Object on page 888 to allow each tenant to customize how DNS resolution is handled within its virtual system. Any service with a **Location** will use the DNS Proxy object configured for the virtual system to determine the primary (or secondary) DNS server to resolve FQDNs, as illustrated in the following figure.

© 2018 Palo Alto Networks, Inc.



STEP 1 | For each virtual system, specify the DNS Proxy to use.

1. Select **Device** > **Virtual Systems** and **Add** the **ID** of the virtual system (range is 1-255), and an optional **Name**, in this example, Corp1 Corporation.
2. On the **General** tab, choose a **DNS Proxy** or create a new one. In this example, Corp1 DNS Proxy is selected as the proxy for Corp1 Corporation's virtual system.
3. For **Interfaces**, click **Add**. In this example, Ethernet1/20 is dedicated to this tenant.
4. For **Virtual Routers**, click **Add**. A virtual router named Corp1 VR is assigned to the virtual system in order to separate routing functions.
5. Click **OK**.

STEP 2 | Configure a DNS Proxy and a server profile to support DNS resolution for a virtual system.

1. Select **Network** > **DNS Proxy** and click **Add**.
2. Click **Enable** and enter a **Name** for the DNS Proxy.
3. For **Location**, select the virtual system of the tenant, in this example, Corp1 Corporation (vsys6). (You could choose the **Shared** DNS Proxy resource instead.)
4. For **Server Profile**, choose or create a profile to customize DNS servers to use for DNS resolutions for this tenant's security policy, reporting, and server profile services.

   If the profile is not already configured, in the **Server Profile** field, click **DNS Server Profile** to Configure a DNS Server Profile on page 891.

   The DNS server profile identifies the IP addresses of the primary and secondary DNS server to use for management DNS resolutions for this virtual system.
5. Also for this server profile, optionally configure a **Service Route IPv4** and/or a **Service Route IPv6** to instruct the firewall which **Source Interface** to use in its DNS requests. If that interface has more than one IP address, configure the **Source Address** also.
6. Click **OK**.
7. Click **OK** and **Commit**.

Optional advanced features such as split DNS can be configured using DNS Proxy Rules. A separate DNS server profile can be used to redirect DNS resolutions matching the Domain Name in a DNS Proxy Rule to another set of DNS servers, if required. Use Case 3 illustrates split DNS.

© 2018 Palo Alto Networks, Inc.

If you use two separate DNS server profiles in the same DNS Proxy object, one for the DNS Proxy and one for the DNS proxy rule, the following behaviors occur:

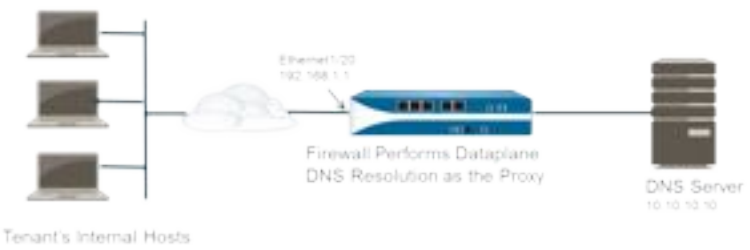
- If a service route is defined in the DNS server profile used by the DNS Proxy, it takes precedence and is used.
- If a service route is defined in the DNS server profile used in the DNS proxy rules, it is not used. If the service route differs from the one defined in the DNS server profile used by the DNS Proxy, the following warning message is displayed during the **Commit** process:

```
Warning: The DNS service route defined in the DNS proxy object is
 different from the DNS proxy rule's service route. Using the DNS proxy
 object's service route.
```

- If no service route is defined in any DNS server profile, the global service route is used if needed.

Use Case 3: Firewall Acts as DNS Proxy Between Client and Server

In this use case, the firewall is located between a DNS client and a DNS server. A DNS Proxy on the firewall is configured to act as the DNS server for the hosts that reside on the tenant's network connected to the firewall interface. In such a scenario, the firewall performs DNS resolution on its dataplane.



This scenario happens to use *split DNS*, a configuration where DNS Proxy rules are configured to redirect DNS requests to a set of DNS servers based on a domain name match. If there is no match, the server profile determines the DNS servers to which to send the request, hence the two, split DNS resolution methods.

For dataplane DNS resolutions, the source IP address from the DNS proxy in PAN-OS to the outside DNS server would be the address of the proxy (the destination IP of the original request). Any service routes defined in the DNS Server Profile are not used. For example, if the request is from host 172.16.1.1 to the DNS proxy at 192.168.1.1, then the request to the DNS server (at 10.10.10.10) would use a source of 192.168.1.1 and a destination of 10.10.10.10.

STEP 1 | Select **Network** > **DNS Proxy** and click **Add**.

STEP 2 | Click **Enable** and enter a **Name** for the DNS Proxy.

STEP 3 | For **Location**, select the virtual system of the tenant, in this example, Corp1 Corporation (vsys6).

STEP 4 | For **Interface**, select the interface that will receive the DNS requests from the tenant's hosts, in this example, Ethernet1/20.

© 2018 Palo Alto Networks, Inc.

STEP 5 | Choose or create a **Server Profile** to customize DNS servers to resolve DNS requests for this tenant.

STEP 6 | On the **DNS Proxy Rules** tab, **Add** a **Name** for the rule.

STEP 7 | (Optional) Select **Turn on caching of domains resolved by this mapping**.

STEP 8 | **Add** one or more **Domain Name**(s), one entry per row. DNS Proxy Rule and FQDN Matching describes how the firewall matches FQDNs to domain names in a DNS proxy rule.

STEP 9 | For **DNS Server profile**, select a profile from the drop-down. The firewall compares the domain name in the DNS request to the domain name(s) defined in the **DNS Proxy Rules**. If there is a match, the **DNS Server profile** defined in the rule is used to determine the DNS server.

STEP 10 | In this example, if the domain in the request matches myweb.corp1.com, the DNS server defined in the myweb DNS Server Profile is used. If there is no match, the DNS server defined in the **Server Profile** (Corp1 DNS Server Profile) is used.

STEP 11 | Click **OK** twice.

DNS Proxy Rule and FQDN Matching

When you configure the firewall with a DNS Proxy Object that uses DNS proxy rules, the firewall compares an FQDN from a DNS query to the domain name of a DNS proxy rule. The firewall comparison works as follows:

FQDN Comparison to DNS Proxy Rule	For Example
The firewall first tokenizes the FQDNs and the domains names in the DNS proxy rules. In a domain name, a string delimited by a period (.) is a token.	**`*.boat.fish.com`** consists of four tokens: **`[*][boat][fish][com]`**
The matching process is an exact token match between the FQDN and the domain name in the rule; partial strings are not matched.	Rule: **`fishing`** FQDN: **`fish`** — **Not a Match**
An exception to the exact match requirement is the use of the wildcard—an asterisk (*). The * matches one or more tokens. This means a rule consisting of only a wildcard (*) matches any FQDN with one or more tokens.	Rule: **`*.boat.com`** FQDN: **`www.boat.com`** — Match FQDN: **`www.blue.boat.com`** — Match FQDN: **`boat.com`** — **Not a Match**
	Rule: **`*`** FQDN: **`boat`** — Match FQDN: **`boat.com`** — Match FQDN: **`www.boat.com`** — Match
You can use an * in any position: preceding tokens, between tokens, or trailing tokens	Rule: **`www.*.com`**

© 2018 Palo Alto Networks, Inc.

FQDN Comparison to DNS Proxy Rule	For Example
(but not with other characters within a single token).	FQDN: **www.boat.com** — Match FQDN: **www.blue.boat.com** — Match
	Rule: **www.boat.*** FQDN: **www.boat.com** — Match FQDN: **www.boat.fish.com** — Match
	Rule: **www.boat*.com** — **Invalid**
Multiple wildcards (*) can appear in any position of the domain name: preceding tokens, between tokens, or trailing tokens. Each non-consecutive * matches one or more tokens.	Rule: **a.*.d.*.com** FQDN: **a.b.d.e.com** — Match FQDN: **a.b.c.d.e.f.com** — Match FQDN: **a.d.d.e.f.com** — Match (First * matches **d**; second * matches **e** and **f**) FQDN: **a.d.e.f.com** — **Not a Match** (First * matches **d**; subsequent **d** in the rule is not matched)
When wildcards are used in consecutive tokens, the first * matches one or more tokens; the second * matches one token. This means a rule consisting of only *.* matches any FQDN with two or more tokens.	Consecutive wildcards preceding tokens: Rule: ***.*.boat.com** FQDN: **www.blue.boat.com** — Match FQDN: **www.blue.sail.boat.com** — Match
	Consecutive wildcards between tokens: Rule: **www.*.*.boat.com** FQDN: **www.blue.sail.boat.com** — Match FQDN: **www.big.blue.sail.boat.com** — Match
	Consecutive wildcards trailing tokens: Rule: **www.boat.*.*** FQDN: **www.boat.fish.com** — Match FQDN: **www.boat.fish.ocean.com** — Match
	Consecutive wildcards only: Rule: *.* FQDN: **boat** — **Not a Match** FQDN: **boat.com** — Match FQDN: **www.boat.com** — Match
Consecutive and non-consecutive wildcards can appear in the same rule.	Rule: **a.*.d.*.*.com** FQDN: **a.b.c.d.e.f.com** — Match (First * matches **b** and **c**; second * matches **e**; third * matches **f**)

© 2018 Palo Alto Networks, Inc.

FQDN Comparison to DNS Proxy Rule	For Example
	FQDN: **a.b.c.d.e.com** — **Not a Match** (First * matches **b** and **c**; second * matches **e**; third * not matched)
The Implicit-tail-match behavior provides an additional shorthand: As long as the last token of the rule is not an *, a comparison will match if all tokens in the rule match the FQDN, even when the FQDN has additional trailing tokens that the rule doesn't have.	Rule: **www.boat.fish** FQDN: **www.boat.fish.com** — Match FQDN: **www.boat.fish.ocean.com** — Match FQDN: **www.boat.fish** — Match
This rule ends with *, so the Implicit-tail-match rule doesn't apply. The * behaves as stated; it matches one or more tokens.	Rule: **www.boat.fish.*** FQDN: **www.boat.fish.com** — Match FQDN: **www.boat.fish.ocean.com** — Match FQDN: **www.boat.fish** — **Not a Match** (This FQDN does not have a token to match the * in the rule.)
In the case where an FQDN matches more than one rule, a tie-breaking algorithm selects the most specific (longest) rule; that is, the algorithm favors the rule with more tokens and fewer wildcards (*).	Rule 1: ***.fish.com** — Match Rule 2: ***.com** — Match Rule 3: **boat.fish.com** — Match and Tie-Breaker FQDN: **boat.fish.com** FQDN matches all three rules; the firewall uses Rule 3 because it is the most specific.
	Rule 1: ***.fish.com** — **Not a Match** Rule 2: ***.com** — Match Rule 3: **boat.fish.com** — **Not a Match** FQDN: **fish.com** FQDN does not match Rule 1 because the * does not have a token to match.
	Rule 1: ***.fish.com** — Match and Tie-Breaker Rule 2: ***.com** — Match Rule 3: **boat.fish.com** — **Not a Match** FQDN: **blue.boat.fish.com** FQDN matches Rule 1 and Rule 2 (because the * matches one or more tokens). The firewall uses Rule 1 because it is the most specific.
When working with wildcards (*) and Implicit-tail-match rules, there can be cases when the FQDN matches more than one rule and the tie-breaking algorithm weighs the rules equally.	Replace this: Rule: **www.boat** with this:

© 2018 Palo Alto Networks, Inc.

FQDN Comparison to DNS Proxy Rule	For Example
To avoid ambiguity, if rules with an Implicit-tail-match or a wildcard (*) can overlap, replace an Implicit-tail-match rule by specifying the tail token.	Rule: **www.boat.com**
Best Practices for Creating DNS Proxy Rules to Avoid Ambiguity and Unexpected Results	
Avoid invoking an Implicit-tail-match by including a top-level domain in the domain name.	**boat.com**
If you use a wildcard (*), use it only as the leftmost token. This practice follows the common understanding of wildcard DNS records and the hierarchical nature of DNS.	***.boat.com**
Use no more than one * in a rule.	
Use the * to establish a base rule associated with a DNS server, and use rules with more tokens to build exceptions to the rule, which you associate with different servers. The tie-breaking algorithm will select the most specific match, based on the number of matched tokens.	Rule: ***.corporation.com** — DNS server A Rule: **www.corporation.com** — DNS server B Rule: ***.internal.corporation.com** — DNS server C Rule: **www.internal.corporation.com** — DNS server D FQDN: **mail.internal.corporation.com** — matches DNS server C FQDN: **mail.corporation.com** — matches DNS server A

© 2018 Palo Alto Networks, Inc.

NAT

This section describes Network Address Translation (NAT) and how to configure the firewall for NAT. NAT allows you to translate private, non-routable IPv4 addresses to one or more globally-routable IPv4 addresses, thereby conserving an organization's routable IP addresses. NAT allows you to not disclose the real IP addresses of hosts that need access to public addresses and to manage traffic by performing port forwarding. You can use NAT to solve network design challenges, enabling networks with identical IP subnets to communicate with each other. The firewall supports NAT on Layer 3 and virtual wire interfaces.

The NAT64 option translates between IPv6 and IPv4 addresses, providing connectivity between networks using disparate IP addressing schemes, and therefore a migration path to IPv6 addressing. IPv6-to-IPv6 Network Prefix Translation (NPTv6) translates one IPv6 prefix to another IPv6 prefix. PAN-OS supports all of these functions.

If you use private IP addresses within your internal networks, you must use NAT to translate the private addresses to public addresses that can be routed on external networks. In PAN-OS, you create NAT policy rules that instruct the firewall which packet addresses and ports need translation and what the translated addresses and ports are.

- NAT Policy Rules
- Source NAT and Destination NAT
- NAT Rule Capacities
- Dynamic IP and Port NAT Oversubscription
- Dataplane NAT Memory Statistics
- Configure NAT
- NAT Configuration Examples

NAT Policy Rules

- NAT Policy Overview on page 900
- NAT Address Pools Identified as Address Objects on page 901
- Proxy ARP for NAT Address Pools on page 902

NAT Policy Overview

You configure a NAT rule to match a packet's source zone and destination zone, at a minimum. In addition to zones, you can configure matching criteria based on the packet's destination interface, source and destination address, and service. You can configure multiple NAT rules. The firewall evaluates the rules in order from the top down. Once a packet matches the criteria of a single NAT rule, the packet is not subjected to additional NAT rules. Therefore, your list of NAT rules should be in order from most specific to least specific so that packets are subjected to the most specific rule you created for them.

Static NAT rules do not have precedence over other forms of NAT. Therefore, for static NAT to work, the static NAT rules must be above all other NAT rules in the list on the firewall.

NAT rules provide address translation, and are different from security policy rules, which allow or deny packets. It is important to understand the firewall's flow logic when it applies NAT rules and security policy rules so that you can determine what rules you need, based on the zones you have defined. You must configure security policy rules to allow the NAT traffic.

Upon ingress, the firewall inspects the packet and does a route lookup to determine the egress interface and zone. Then the firewall determines if the packet matches one of the NAT rules that have been defined, based on source and/or destination zone. It then evaluates and applies any security policies that match the packet based on the original (pre-NAT) source and destination addresses, but the post-NAT zones. Finally,

© 2018 Palo Alto Networks, Inc.

upon egress, for a matching NAT rule, the firewall translates the source and/or destination address and port numbers.

Keep in mind that the translation of the IP address and port do not occur until the packet leaves the firewall. The NAT rules and security policies apply to the original IP address (the pre-NAT address). A NAT rule is configured based on the zone associated with a pre-NAT IP address.

Security policies differ from NAT rules because security policies examine post-NAT zones to determine whether the packet is allowed or not. Because the very nature of NAT is to modify source or destination IP addresses, which can result in modifying the packet's outgoing interface and zone, security policies are enforced on the post-NAT zone.



A SIP call sometimes experiences one-way audio when going through the firewall because the call manager sends a SIP message on behalf of the phone to set up the connection. When the message from the call manager reaches the firewall, the SIP ALG must put the IP address of the phone through NAT. If the call manager and the phones are not in the same security zone, the NAT lookup of the IP address of the phone is done using the call manager zone. The NAT policy should take this into consideration.

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

You can verify the NAT rules processed by using the CLI **`test nat-policy-match`** command in operational mode. For example:

```
user@device1> test nat-policy-match
?
+ destination  Destination IP address
+ destination-port  Destination port
+ from  From zone
+ ha-device-id  HA Active/Active device ID
+ protocol  IP protocol value
+ source  Source IP address
+ source-port  Source port
+ to  To Zone
+ to-interface  Egress interface to use
  |  Pipe through a command
<Enter>  Finish input
user@device1> test nat-policy-match from l3-untrust source
10.1.1.1 destination 66.151.149.20 destination-port 443 protocol
6
Destination-NAT: Rule matched: CA2-DEMO
66.151.149.20:443 => 192.168.100.15:443
```

NAT Address Pools Identified as Address Objects

When configuring a **Dynamic IP** or **Dynamic IP and Port** NAT address pool in a NAT policy rule, it is typical to configure the pool of translated addresses with address objects. Each address object can be a host IP address, IP address range, or IP subnet.



Because both NAT rules and security policy rules use address objects, it is a best practice to distinguish between them by naming an address object used for NAT with a prefix, such as "NAT-name."

© 2018 Palo Alto Networks, Inc.

Proxy ARP for NAT Address Pools

NAT address pools are not bound to any interfaces. The following figure illustrates the behavior of the firewall when it is performing proxy ARP for an address in a NAT address pool.



The firewall performs source NAT for a client, translating the source address 10.1.1.1 to the address in the NAT pool, 192.168.2.2. The translated packet is sent on to a router.

For the return traffic, the router does not know how to reach 192.168.2.2 (because that IP address is just an address in the NAT address pool), so it sends an ARP request packet to the firewall.

- If the address pool (192.168.2.2) is in the same subnet as the egress/ingress interface IP address (192.168.2.3/24), the firewall can send a proxy ARP reply to the router, indicating the Layer 2 MAC address of the IP address, as shown in the figure above.
- If the address pool (192.168.2.2) is not a subnet of an interface on the firewall, the firewall will not send a proxy ARP reply to the router. This means that the router must be configured with the necessary route to know where to send packets destined for 192.168.2.2, in order to ensure the return traffic is routed back to the firewall, as shown in the figure below.



Source NAT and Destination NAT

The firewall supports both source address and/or port translation and destination address and/or port translation.

- Source NAT on page 902
- Destination NAT on page 903

Source NAT

Source NAT is typically used by internal users to access the Internet; the source address is translated and thereby kept private. There are three types of source NAT:

- **Dynamic IP and Port (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 will be 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 (NAPT).

  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. For more information, see Dynamic IP and Port NAT Oversubscription on page 904 and Modify the Oversubscription Rate for DIPP NAT on page 910.

© 2018 Palo Alto Networks, Inc.

- **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 on page 911.
- **Static IP**—Allows the 1-to-1, 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.

Destination NAT

Destination NAT is performed on incoming packets, when the firewall translates a public destination address to a private address. Destination NAT also offers the option to perform port forwarding or port translation.

Destination NAT is a one-to-one, static translation that you can configure in several formats. You can specify that the original packet have a single destination IP address, a range of IP addresses, or a list of single IP addresses, as long as the translated packet is in the same format and specifies the same number of IP addresses. The firewall statically translates an original destination address to the same translated destination address each time. That is, if there is more than one destination address, the firewall translates the first destination address configured for the original packet to the first destination address configured for the translated packet, and translates the second original destination address configured to the second translated destination address configured, and so on, always using the same translation.

For example, the firewall allows the following destination NAT translations:

Original Packet's Destination Address	Maps to Translated Packet's Destination Address	Notes
192.168.1.1	2.2.2.2	Original packet and translated packet each have one possible destination address.
192.168.1.1-192.168.1.4	2.2.2.1-2.2.2.4	Original packet and translated packet each have four possible destination addresses: 192.168.1.1 always maps to 2.2.2.1 192.168.1.2 always maps to 2.2.2.2 192.168.1.3 always maps to 2.2.2.3 192.168.1.4 always maps to 2.2.2.4
192.168.1.7 192.168.1.4 192.168.1.253 192.168.1.1	2.2.2.1 2.2.2.2 2.2.2.3 2.2.2.4	Original packet and translated packet each have four possible destination addresses: 192.168.1.7 always maps to 2.2.2.1 192.168.1.4 always maps to 2.2.2.2 192.168.1.253 always maps to 2.2.2.3

© 2018 Palo Alto Networks, Inc.

Original Packet's Destination Address	Maps to Translated Packet's Destination Address	Notes
		192.168.1.1 always maps to 2.2.2.4
192.168.1.1/30	2.2.2.1/30	Original packet and translated packet each have four possible destination addresses: 192.168.1.1 always maps to 2.2.2.1 192.168.1.2 always maps to 2.2.2.2 192.168.1.3 always maps to 2.2.2.3 192.168.1.4 always maps to 2.2.2.4

One common use of destination NAT is to configure several NAT rules that map a single public destination address to several private destination host addresses assigned to servers or services. In this case, the destination port numbers are used to identify the destination hosts. For example:

- **Port Forwarding**—Can translate a public destination address and port number to a private destination address, but keeps the same port number.
- **Port Translation**—Can translate a public destination address and port number to a private destination address and a different port number, thus keeping the real port number private. It is configured by entering a **Translated Port** on the **Translated Packet** tab in the NAT policy rule. See the Destination NAT with Port Translation Example.

NAT Rule Capacities

The number of NAT rules allowed is based on the firewall model. Individual rule limits are set for static, Dynamic IP (DIP), and Dynamic IP and Port (DIPP) NAT. The sum of the number of rules used for these NAT types cannot exceed the total NAT rule capacity. For DIPP, the rule limit is based on the oversubscription setting (8, 4, 2, or 1) of the firewall and the assumption of one translated IP address per rule. To see model-specific NAT rule limits and translated IP address limits, use the Compare Firewalls tool.

Consider the following when working with NAT rules:

- If you run out of pool resources, you cannot create more NAT rules, even if the model's maximum rule count has not been reached.
- If you consolidate NAT rules, the logging and reporting will also be consolidated. The statistics are provided per the rule, not per all of the addresses within the rule. If you need granular logging and reporting, do not combine the rules.

Dynamic IP and Port NAT Oversubscription

Dynamic IP and Port (DIPP) NAT allows you to use each translated IP address and port pair multiple times (8, 4, or 2 times) in concurrent sessions. This reusability of an IP address and port (known as oversubscription) provides scalability for customers who have too few public IP addresses. The design is based on the assumption that hosts are connecting to different destinations, therefore sessions can be uniquely identified and collisions are unlikely. The oversubscription rate in effect multiplies the original size of the address/port pool to 8, 4, or 2 times the size. For example, the default limit of 64K concurrent sessions allowed, when multiplied by an oversubscription rate of 8, results in 512K concurrent sessions allowed.

© 2018 Palo Alto Networks, Inc.

The oversubscription rates that are allowed vary based on the model. The oversubscription rate is global; it applies to the firewall. This oversubscription rate is set by default and consumes memory, even if you have enough public IP addresses available to make oversubscription unnecessary. You can reduce the rate from the default setting to a lower setting or even 1 (which means no oversubscription). By configuring a reduced rate, you decrease the number of source device translations possible, but increase the DIP and DIPP NAT rule capacities. To change the default rate, see Modify the Oversubscription Rate for DIPP NAT on page 910.

If you select **Platform Default**, your explicit configuration of oversubscription is turned off and the default oversubscription rate for the model applies, as shown in the table below. The **Platform Default** setting allows for an upgrade or downgrade of a software release.

The following table lists the default (highest) oversubscription rate for each model.

Model	Default Oversubscription Rate
PA-200	2
PA-220	2
PA-500	2
PA-820	2
PA-850	2
PA-3020	2
PA-3050	2
PA-3060	2
PA-5020	4
PA-5050	8
PA-5060	8
PA-5220	4
PA-5250	8
PA-5260	8
PA-7050	8
PA-7080	8
VM-50	2
VM-100	1
VM-200	1

© 2018 Palo Alto Networks, Inc.

Model	Default Oversubscription Rate
VM-300	2
VM-500	8
VM-700	8
VM-1000-HV	2

The firewall supports a maximum of 256 translated IP addresses per NAT rule, and each model supports a maximum number of translated IP addresses (for all NAT rules combined). If oversubscription causes the maximum translated addresses per rule (256) to be exceeded, the firewall will automatically reduce the oversubscription ratio in an effort to have the commit succeed. However, if your NAT rules result in translations that exceed the maximum translated addresses for the model, the commit will fail.

Dataplane NAT Memory Statistics

The **`show running global-ippool`** command displays statistics related to NAT memory consumption for a pool. The Size column displays the number of bytes of memory that the resource pool is using. The Ratio column displays the oversubscription ratio (for DIPP pools only). The lines of pool and memory statistics are explained in the following sample output:



For NAT pool statistics for a virtual system, the **`show running ippool`** command has columns indicating the memory size used per NAT rule and the oversubscription ratio used (for DIPP rules). The following is sample output for the command.



A field in the output of the **`show running nat-rule-ippool rule`** command shows the memory (bytes) used per NAT rule. The following is sample output for the command, with the memory usage for the rule encircled.

© 2018 Palo Alto Networks, Inc.

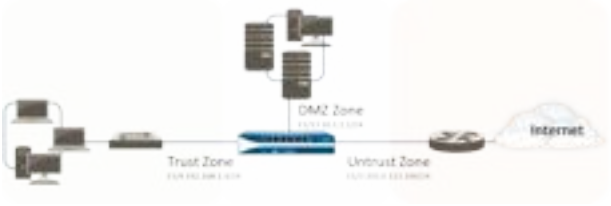


Configure NAT

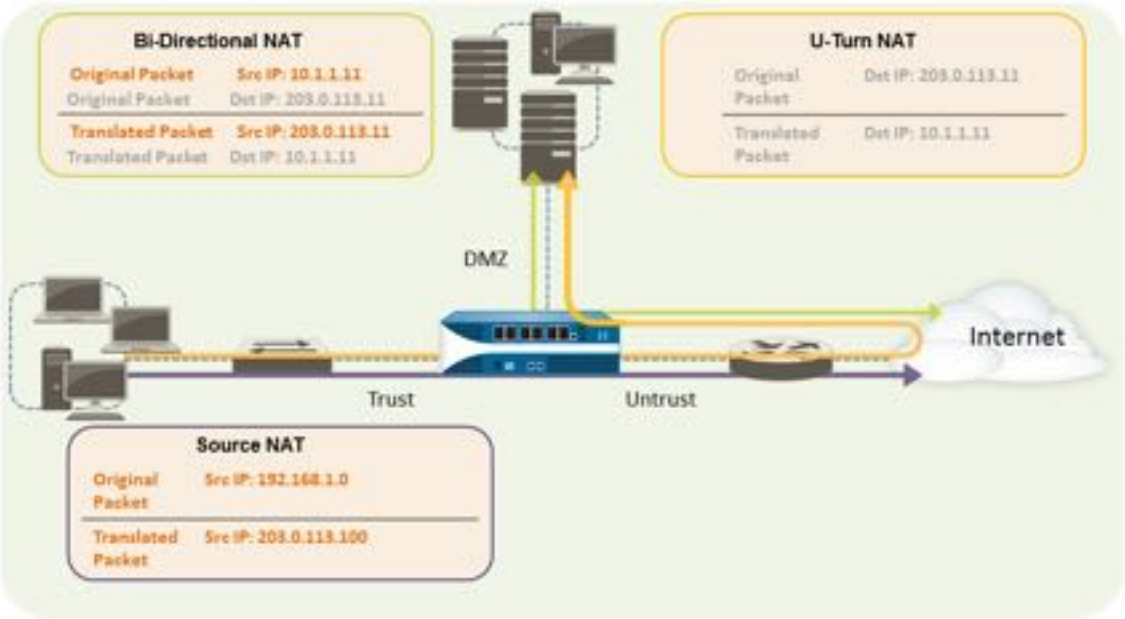
Perform the following tasks to configure various aspects of NAT. In addition to the examples below, there are examples in the section NAT Configuration Examples.

- Translate Internal Client IP Addresses to Your Public IP Address (Source DIPP NAT)
- Enable Clients on the Internal Network to Access your Public Servers (Destination U-Turn NAT)
- Enable Bi-Directional Address Translation for Your Public-Facing Servers (Static Source NAT)
- Modify the Oversubscription Rate for DIPP NAT
- Disable NAT for a Specific Host or Interface
- Reserve Dynamic IP NAT Addresses

The NAT example in this section is based on the following topology:



Based on this topology, there are three NAT policies we need to create as follows:



© 2018 Palo Alto Networks, Inc.

- To enable the clients on the internal network to access resources on the Internet, the internal 192.168.1.0 addresses will need to be translated to publicly routable addresses. In this case, we will configure source NAT (the purple enclosure and arrow above), using the egress interface address, 203.0.113.100, as the source address in all packets that leave the firewall from the internal zone. See Translate Internal Client IP Addresses to Your Public IP Address (Source DIPP NAT) for instructions.
- To enable clients on the internal network to access the public web server in the DMZ zone, we must configure a NAT rule that redirects the packet from the external network, where the original routing table lookup will determine it should go based on the destination address of 203.0.113.11 within the packet, to the actual address of the web server on the DMZ network of 10.1.1.11. To do this you must create a NAT rule from the trust zone (where the source address in the packet is) to the untrust zone (where the original destination address is) to translate the destination address to an address in the DMZ zone. This type of destination NAT is called *U-Turn NAT* (the yellow enclosure and arrow above). See Enable Clients on the Internal Network to Access your Public Servers (Destination U-Turn NAT) for instructions.
- To enable the web server—which has both a private IP address on the DMZ network and a public-facing address for access by external users—to both send and receive requests, the firewall must translate the incoming packets from the public IP address to the private IP address and the outgoing packets from the private IP address to the public IP address. On the firewall, you can accomplish this with a single bi-directional static source NAT policy (the green enclosure and arrow above). See Enable Bi-Directional Address Translation for Your Public-Facing Servers (Static Source NAT).

Translate Internal Client IP Addresses to Your Public IP Address (Source DIPP NAT)

When a client on your internal network sends a request, the source address in the packet contains the IP address for the client on your internal network. If you use private IP address ranges internally, the packets from the client will not be able to be routed on the Internet unless you translate the source IP address in the packets leaving the network into a publicly routable address.

On the firewall you can do this by configuring a source NAT policy that translates the source address (and optionally the port) into a public address. One way to do this is to translate the source address for all packets to the egress interface on your firewall, as shown in the following procedure.

STEP 1 | Create an address object for the external IP address you plan to use.

1. Select **Objects** > **Addresses** and **Add** a **Name** and optional **Description** for the object.
2. Select **IP Netmask** from the **Type** drop-down and then enter the IP address of the external interface on the firewall, 203.0.113.100 in this example.
3. Click **OK**.



Although you do not have to use address objects in your policies, it is a best practice because it simplifies administration by allowing you to make updates in one place rather than having to update every policy where the address is referenced.

STEP 2 | Create the NAT policy.

1. Select **Policies** > **NAT** and click **Add**.
2. On the **General** tab, enter a descriptive **Name** for the policy.
3. (Optional) Enter a tag, which is a keyword or phrase that allows you to sort or filter policies.
4. For **NAT Type**, select **ipv4** (default).
5. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.
6. On the **Translated Packet** tab, select **Dynamic IP And Port** from the **Translation Type** drop-down in the Source Address Translation section of the screen.

© 2018 Palo Alto Networks, Inc.

7. For **Address Type**, there are two choices. You could select **Translated Address** and then click **Add**. Select the address object you just created.

   An alternative **Address Type** is **Interface Address**, in which case the translated address will be the IP address of the interface. For this choice, you would select an **Interface** and optionally an **IP Address** if the interface has more than one IP address.
8. Click **OK**.

STEP 3 | Commit your changes.

Click **Commit**.

STEP 4 | (Optional) Access the CLI to verify the translation.

1. Use the **`show session all`** command to view the session table, where you can verify the source IP address and port and the corresponding translated IP address and port.
2. Use the *show session id <id_number>* to view more details about a session.
3. If you configured Dynamic IP NAT, use the **`show counter global filter aspect session severity drop | match nat`** command to see if any sessions failed due to NAT IP allocation. If all of the addresses in the Dynamic IP NAT pool are allocated when a new connection is supposed to be translated, the packet will be dropped.

Enable Clients on the Internal Network to Access your Public Servers (Destination U-Turn NAT)

When a user on the internal network sends a request for access to the corporate web server in the DMZ, the DNS server will resolve it to the public IP address. When processing the request, the firewall will use the original destination in the packet (the public IP address) and route the packet to the egress interface for the untrust zone. In order for the firewall to know that it must translate the public IP address of the web server to an address on the DMZ network when it receives requests from users on the trust zone, you must create a destination NAT rule that will enable the firewall to send the request to the egress interface for the DMZ zone as follows.

STEP 1 | Create an address object for the web server.

1. Select **Objects** > **Addresses** and **Add** a **Name** and optional **Description** for the object.
2. Select **IP Netmask** from the **Type** drop-down and enter the public IP address of the web server, 203.0.113.11 in this example.
3. Click **OK**.

STEP 2 | Create the NAT policy.

1. Select **Policies** > **NAT** and click **Add**.
2. On the **General** tab, enter a descriptive **Name** for the NAT rule.
3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.
4. In the **Destination Address** section, **Add** the address object you created for your public web server.
5. On the **Translated Packet** tab, select **Destination Address Translation** and then enter the IP address that is assigned to the web server interface on the DMZ network, 10.1.1.11 in this example.
6. Click **OK**.

STEP 3 | Commit.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Enable Bi-Directional Address Translation for Your Public-Facing Servers (Static Source NAT)

When your public-facing servers have private IP addresses assigned on the network segment where they are physically located, you need a source NAT rule to translate the source address of the server to the external address upon egress. You create a static NAT rule to translate the internal source address, 10.1.1.11, to the external web server address, 203.0.113.11 in our example.

However, a public-facing server must be able to both send and receive packets. You need a reciprocal policy that translates the public address (the destination IP address in incoming packets from Internet users) into the private address so that the firewall can route the packet to your DMZ network. You create a bi-directional static NAT rule, as described in the following procedure. Bi-directional translation is an option for static NAT only.

STEP 1 | Create an address object for the web server's internal IP address.

1. Select **Objects** > **Addresses** and **Add** a **Name** and optional **Description** for the object.
2. Select **IP Netmask** from the **Type** drop-down and enter the IP address of the web server on the DMZ network, 10.1.1.11 in this example.
3. Click **OK**.



If you did not already create an address object for the public address of your web server, you should create that object now.

STEP 2 | Create the NAT policy.

1. Select **Policies** > **NAT** and click **Add**.
2. On the **General** tab, enter a descriptive **Name** for the NAT rule.
3. On the **Original Packet** tab, select the zone you created for your DMZ in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.
4. In the **Source Address** section, **Add** the address object you created for your internal web server address.
5. On the **Translated Packet** tab, select **Static IP** from the **Translation Type** drop-down in the **Source Address Translation** section and then select the address object you created for your external web server address from the **Translated Address** drop-down.
6. In the **Bi-directional** field, select **Yes**.
7. Click **OK**.

STEP 3 | Commit.

Click **Commit**.

Modify the Oversubscription Rate for DIPP NAT

If you have enough public IP addresses that you do not need to use DIPP NAT oversubscription, you can reduce the oversubscription rate and thereby gain more DIP and DIPP NAT rules allowed.

STEP 1 | View the DIPP NAT oversubscription rate.

1. Select **Device** > **Setup** > **Session** > **Session Settings**. View the **NAT Oversubscription Rate** setting.

STEP 2 | Set the DIPP NAT oversubscription rate.

1. Edit the Session Settings section.
2. In the **NAT Oversubscription Rate** drop-down, select **1x**, **2x**, **4x**, or **8x**, depending on which ratio you want.

© 2018 Palo Alto Networks, Inc.

The Platform Default setting applies the default oversubscription setting for the model. If you want no oversubscription, select 1x.

3. Click **OK** and **Commit** the change.

Disable NAT for a Specific Host or Interface

Both source NAT and destination NAT rules can be configured to disable address translation. You may have exceptions where you do not want NAT to occur for a certain host in a subnet or for traffic exiting a specific interface. The following procedure shows how to disable source NAT for a host.

STEP 1 | Create the NAT policy.

1. Select **Policies** > **NAT** and click **Add** a descriptive **Name** for the policy.
2. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.
3. For **Source Address**, click **Add** and enter the host address. Click **OK**.
4. On the **Translated Packet** tab, select **None** from the **Translation Type** drop-down in the Source Address Translation section of the screen.
5. Click **OK**.

STEP 2 | Commit your changes.

Click **Commit**.

NAT rules are processed in order from the top to the bottom, so place the NAT exemption policy before other NAT policies to ensure it is processed before an address translation occurs for the sources you want to exempt.

Reserve Dynamic IP NAT Addresses

You can reserve Dynamic IP NAT addresses (for a configurable period of time) to prevent them from being allocated as translated addresses to a different source IP address that needs translation. When configured, the reservation applies to all of the translated Dynamic IP addresses in progress and any new translations.

For both translations in progress and new translations, when a source IP address is translated to an available translated IP address, that pairing is retained even after all sessions related to that specific source IP are expired. The reservation timer for each source IP address begins after all sessions that use that source IP address translation expire. Dynamic IP NAT is a one-to-one translation; one source IP address translates to one translated IP address that is chosen dynamically from those addresses available in the configured pool. Therefore, a translated IP address that is reserved is not available for any other source IP address until the reservation expires because a new session has not started. The timer is reset each time a new session for a source IP/translated IP mapping begins, after a period when no sessions were active.

By default, no addresses are reserved. You can reserve Dynamic IP NAT addresses for the firewall or for a virtual system.

- Reserve dynamic IP NAT addresses for a firewall.

  Enter the following commands:

```
admin@PA-3020# set setting nat reserve-ip yes
```


© 2018 Palo Alto Networks, Inc.

```
admin@PA-3020# set setting nat reserve-time <1-604800 secs>
```

- Reserve dynamic IP NAT addresses for a virtual system.

  Enter the following commands:

```
admin@PA-3020# set vsys <vsysid> setting nat reserve-ip yes
```

```
admin@PA-3020# set vsys <vsysid> setting nat reserve-time <1-604800 secs>
```

  For example, suppose there is a Dynamic IP NAT pool of 30 addresses and there are 20 translations in progress when the **nat reserve-time** is set to 28800 seconds (8 hours). Those 20 translations are now reserved, so that when the last session (of any application) that uses each source IP/translated IP mapping expires, the translated IP address is reserved for only that source IP address for 8 hours, in case that source IP address needs translation again. Additionally, as the 10 remaining translated addresses are allocated, they each are reserved for their source IP address, each with a timer that begins when the last session for that source IP address expires.

  In this manner, each source IP address can be repeatedly translated to its same NAT address from the pool; another host will not be assigned a reserved translated IP address from the pool, even if there are no active sessions for that translated address.

  Suppose a source IP/translated IP mapping has all of its sessions expire, and the reservation timer of 8 hours begins. After a new session for that translation begins, the timer stops, and the sessions continue until they all end, at which point the reservation timer starts again, reserving the translated address.

  The reservation timer remain in effect on the Dynamic IP NAT pool until you disable it by entering the **set setting nat reserve-ip no** command or you change the **nat reserve-time** to a different value.

  The CLI commands for reservations do not affect Dynamic IP and Port (DIPP) or Static IP NAT pools.

NAT Configuration Examples

- Destination NAT Example—One-to-One Mapping on page 912
- Destination NAT with Port Translation Example on page 914
- Destination NAT Example—One-to-Many Mapping on page 914
- Source and Destination NAT Example on page 915
- Virtual Wire Source NAT Example on page 916
- Virtual Wire Static NAT Example on page 917
- Virtual Wire Destination NAT Example on page 917

Destination NAT Example—One-to-One Mapping

The most common mistakes when configuring NAT and security rules are the references to the zones and address objects. The addresses used in destination NAT rules always refer to the original IP address in the packet (that is, the pre-translated address). The destination zone in the NAT rule is determined after the route lookup of the destination IP address in the original packet (that is, the pre-NAT destination IP address).

The addresses in the security policy also refer to the IP address in the original packet (that is, the pre-NAT address). However, the destination zone is the zone where the end host is physically connected. In

© 2018 Palo Alto Networks, Inc.

other words, the destination zone in the security rule is determined after the route lookup of the post-NAT destination IP address.

In the following example of a one-to-one destination NAT mapping, users from the zone named Untrust-L3 access the server 10.1.1.100 in the zone named DMZ using the IP address 192.0.2.100.



Before configuring the NAT rules, consider the sequence of events for this scenario.

- Host 192.0.2.250 sends an ARP request for the address 192.0.2.100 (the public address of the destination server).
- The firewall receives the ARP request packet for destination 192.0.2.100 on the Ethernet1/1 interface and processes the request. The firewall responds to the ARP request with its own MAC address because of the destination NAT rule configured.
- The NAT rules are evaluated for a match. For the destination IP address to be translated, a destination NAT rule from zone Untrust-L3 to zone Untrust-L3 must be created to translate the destination IP of 192.0.2.100 to 10.1.1.100.
- After determining the translated address, the firewall performs a route lookup for destination 10.1.1.100 to determine the egress interface. In this example, the egress interface is Ethernet1/2 in zone DMZ.
- The firewall performs a security policy lookup to see if the traffic is permitted from zone Untrust-L3 to DMZ.

  The direction of the policy matches the ingress zone and the zone where the server is physically located.

  The security policy refers to the IP address in the original packet, which has a destination address of 192.0.2.100.
- The firewall forwards the packet to the server out egress interface Ethernet1/2. The destination address is changed to 10.1.1.100 as the packet leaves the firewall.

For this example, address objects are configured for webserver-private (10.1.1.100) and Webserver-public (192.0.2.100). The configured NAT rule would look like this:

	Original Packet						Translated Packet	
Name	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
Dst NAT-webserver	Untrust-L3	Untrust-L3	any	any	Webserver-public	any	none	address: webserver-private

The direction of the NAT rules is based on the result of route lookup.

The configured security policy to provide access to the server from the Untrust-L3 zone would look like this:

	Source		Destination						
Name	Zone	Address	Zone	Address	Application	Service	Action	Profile	Options
Webserver access	Untrust-L3	any	DMZ	Webserver-public	web-browsing	any	Allow	none	

© 2018 Palo Alto Networks, Inc.

Destination NAT with Port Translation Example

In this example, the web server is configured to listen for HTTP traffic on port 8080. The clients access the web server using the IP address 192.0.2.100 and TCP Port 80. The destination NAT rule is configured to translate both IP address and port to 10.1.1.100 and TCP port 8080. Address objects are configured for webserver-private (10.1.1.100) and Servers-public (192.0.2.100).



The following NAT and security rules must be configured on the firewall:

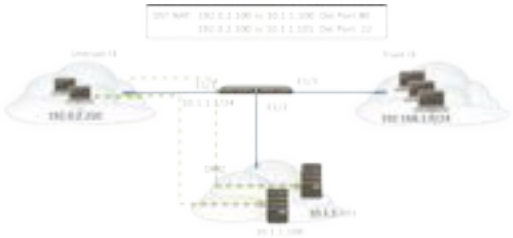


Name	Source Zone	Source Address	Destination Zone	Destination Address	Application	Service
Webserver access	Untrust-L3	any	DMZ	Servers-public	web-browsing	any

Use the **`show session all`** CLI command to verify the translation.

Destination NAT Example—One-to-Many Mapping

In this example, one IP address maps to two different internal hosts. The firewall uses the application to identify the internal host to which the firewall forwards the traffic.



All HTTP traffic is sent to host 10.1.1.100 and SSH traffic is sent to server 10.1.1.101. The following address objects are required:

- Address object for the one pre-translated IP address of the server
- Address object for the real IP address of the SSH server
- Address object for the real IP address of the web server

The corresponding address objects are created:

- Servers-public: 192.0.2.100
- SSH-server: 10.1.1.101
- webserver-private: 10.1.1.100

The NAT rules would look like this:



© 2018 Palo Alto Networks, Inc.

The security rules would look like this:



Source and Destination NAT Example

In this example, NAT rules translate both the source and destination IP address of packets between the clients and the server.

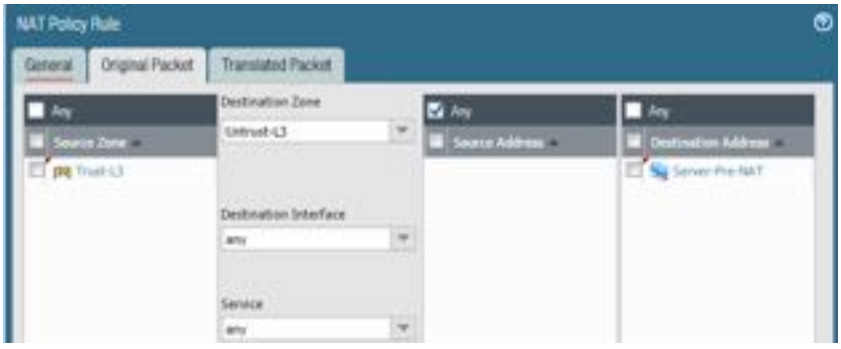
- Source NAT—The source addresses in the packets from the clients in the Trust-L3 zone to the server in the Untrust-L3 zone are translated from the private addresses in the network 192.168.1.0/24 to the IP address of the egress interface on the firewall (10.16.1.103). Dynamic IP and Port translation causes the port numbers to be translated also.
- Destination NAT—The destination addresses in the packets from the clients to the server are translated from the server's public address (80.80.80.80) to the server's private address (10.2.133.15).

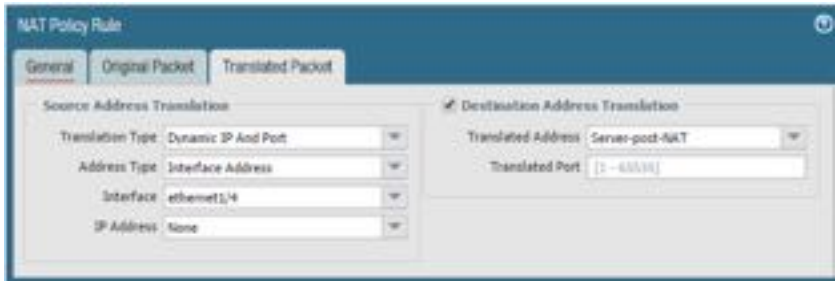


The following address objects are created for destination NAT.

- Server-Pre-NAT: 80.80.80.80
- Server-post-NAT: 10.2.133.15

The following screen shots illustrate how to configure the source and destination NAT policies for the example.





To verify the translations, use the CLI command **`show session all filter destination 80.80.80.80`**. A client address 192.168.1.11 and its port number are translated to 10.16.1.103 and a port number. The destination address 80.80.80.80 is translated to 10.2.133.15.

© 2018 Palo Alto Networks, Inc.

Virtual Wire Source NAT Example

Virtual wire deployment of a Palo Alto Networks firewall includes the benefit of providing security transparently to the end devices. It is possible to configure NAT for interfaces configured in a virtual wire. All of the NAT types are allowed: source NAT (Dynamic IP, Dynamic IP and Port, static) and destination NAT.

Because interfaces in a virtual wire do not have an IP address assigned, it is not possible to translate an IP address to an interface IP address. You must configure an IP address pool.

When performing NAT on virtual wire interfaces, it is recommended that you translate the source address to a different subnet than the one on which the neighboring devices are communicating. The firewall will not proxy ARP for NAT addresses. Proper routing must be configured on the upstream and downstream routers in order for the packets to be translated in virtual wire mode. Neighboring devices will only be able to resolve ARP requests for IP addresses that reside on the interface of the device on the other end of the virtual wire. See Proxy ARP for NAT Address Pools for more explanation about proxy ARP.

In the source NAT example below, security policies (not shown) are configured from the virtual wire zone named vw-trust to the zone named vw-untrust.

In the following topology, two routers are configured to provide connectivity between subnets 192.0.2.0/24 and 172.16.1.0/24. The link between the routers is configured in subnet 198.51.100.0/30. Static routing is configured on both routers to establish connectivity between the networks. Before the firewall is deployed in the environment, the topology and the routing table for each router look like this:



Route on R1:

Destination	Next Hop
172.16.1.0/24	198.51.100.2

Route on R2:

Destination	Next Hop
192.0.2.0/24	198.51.100.1

Now the firewall is deployed in virtual wire mode between the two Layer 3 devices. A NAT IP address pool with range 198.51.100.9 to 198.51.100.14 is configured on the firewall. All communications from clients in network 192.0.2.0/24 accessing servers in network 172.16.1.0/24 will arrive at R2 with a translated source address in the range 198.51.100.9 to 198.51.100.14. The response from servers will be directed to these addresses.



In order for source NAT to work, you must configure proper routing on R2, so that packets destined for other addresses are not dropped. The routing table below shows the modified routing table on R2; the

© 2018 Palo Alto Networks, Inc.

route ensures the traffic to the destinations 198.51.100.9 to 198.51.100.14 (that is, hosts on subnet 198.51.100.8/29) will be sent back through the firewall to router R1.

Route on R2:

Destination	Next Hop
198.51.100.8/29	198.51.100.1

Virtual Wire Static NAT Example

In this example, security policies are configured from the virtual wire zone named Trust to the virtual wire zone named Untrust. Host 192.0.2.100 is statically translated to address 198.51.100.100. With the **Bi-directional** option enabled, the firewall generates a NAT policy from the Untrust zone to the Trust zone. Clients on the Untrust zone access the server using the IP address 198.51.100.100, which the firewall translates to 192.0.2.100. Any connections initiated by the server at 192.0.2.100 are translated to source IP address 198.51.100.100.



Route on R2:

Destination	Next Hop
198.51.100.100/32	198.51.100.1



Virtual Wire Destination NAT Example

Clients in the Untrust zone access the server using the IP address 198.51.100.100, which the firewall translates to 192.0.2.100. Both the NAT and security policies must be configured from the Untrust zone to the Trust zone.



Route on R2:

Destination	Next Hop
198.51.100.100/32	198.51.100.1

© 2018 Palo Alto Networks, Inc.



© 2018 Palo Alto Networks, Inc.

NPTv6

IPv6-to-IPv6 Network Prefix Translation (NPTv6) performs a stateless, static translation of one IPv6 prefix to another IPv6 prefix (port numbers are not changed). There are four primary benefits of NPTv6:

- You can prevent the asymmetrical routing problems that result from Provider Independent addresses being advertised from multiple datacenters.
- NPTv6 allows more specific routes to be advertised so that return traffic arrives at the same firewall that transmitted the traffic.
- Private and public addresses are independent; you can change one without affecting the other.
- You have the ability to translate Unique Local Addresses on page 920 to globally routable addresses.

This topic builds on a basic understanding of NAT. You should be sure you are familiar with NAT on page 900 concepts before configuring NPTv6.

- NPTv6 Overview on page 919
- How NPTv6 Works on page 920
- NDP Proxy on page 922
- NPTv6 and NDP Proxy Example on page 923
- Create an NPTv6 Policy on page 924

NPTv6 Overview

This section describes IPv6-to-IPv6 Network Prefix Translation (NPTv6) and how to configure it. NPTv6 is defined in RFC 6296. Palo Alto Networks does not implement all functionality defined in the RFC, but is compliant with the RFC in the functionality it has implemented.

NPTv6 performs stateless translation of one IPv6 prefix to another IPv6 prefix. It is stateless, meaning that it does not keep track of ports or sessions on the addresses translated. NPTv6 differs from NAT66, which is stateful. Palo Alto Networks supports NPTv6 RFC 6296 prefix translation; it does not support NAT66.

With the limited addresses in the IPv4 space, NAT on page 900 was required to translate private, non-routable IPv4 addresses to one or more globally-routable IPv4 addresses.

For organizations using IPv6 addressing, there is no need to translate IPv6 addresses to IPv6 addresses due to the abundance of IPv6 addresses. However, there are Reasons to Use NPTv6 on page 920 to translate IPv6 prefixes at the firewall.

NPTv6 translates the prefix portion of an IPv6 address but not the host portion or the application port numbers. The host portion is simply copied, and therefore remains the same on either side of the firewall. The host portion also remains visible within the packet header.

- NPTv6 Does Not Provide Security on page 919
- Model Support for NPTv6 on page 920
- Unique Local Addresses on page 920
- Reasons to Use NPTv6 on page 920

NPTv6 Does Not Provide Security

It is important to understand that NPTv6 does not provide security. In general, stateless network address translation does not provide any security; it provides an address translation function. NPTv6 does not hide or translate port numbers. You must set up firewall security policies correctly in each direction to ensure that traffic is controlled as you intended.

© 2018 Palo Alto Networks, Inc.

Model Support for NPTv6

NPTv6 is supported on the following models (NPTv6 with hardware lookup but packets go through the CPU): PA-7000 Series, PA-5200 Series, PA-5000 Series, PA-3060 firewall, and PA-3050 firewall, PA-800 firewall and PA-220 firewall. Models supported with no ability to have hardware perform a session look-up: PA-3020 firewall, PA-500 firewall, PA-200 firewall, and VM-Series.

Unique Local Addresses

RFC 4193, Unique Local IPv6 Unicast Addresses, defines unique local addresses (ULAs), which are IPv6 unicast addresses. They can be considered IPv6 equivalents of the private IPv4 addresses identified in RFC 1918, Address Allocation for Private Internets, which cannot be routed globally.

A ULA is globally unique, but not expected to be globally routable. It is intended for local communications and to be routable in a limited area such as a site or among a small number of sites. Palo Alto Networks does not recommend that you assign ULAs, but a firewall configured with NPTv6 will translate prefixes sent to it, including ULAs.

Reasons to Use NPTv6

Although there is no shortage of public, globally routable IPv6 addresses, there are reasons you might want to translate IPv6 addresses. NPTv6:

- **Prevents asymmetrical routing**—Asymmetric routing can occur if a Provider Independent address space (/48, for example) is advertised by multiple data centers to the global Internet. By using NPTv6, you can advertise more specific routes from regional firewalls, and the return traffic will arrive at the same firewall where the source IP address was translated by the translator.
- **Provides address independence**—You need not change the IPv6 prefixes used inside your local network if the global prefixes are changed (for example, by an ISP or as a result of merging organizations). Conversely, you can change the inside addresses at will without disrupting the addresses that are used to access services in the private network from the Internet. In either case, you update a NAT rule rather than reassign network addresses.
- **Translates ULAs for routing**—You can have Unique Local Addresses on page 920 assigned within your private network, and have the firewall translate them to globally routable addresses. Thus, you have the convenience of private addressing and the functionality of translated, routable addresses.
- **Reduces exposure to IPv6 prefixes**—IPv6 prefixes are less exposed than if you didn't translate network prefixes, however, NPTv6 is not a security measure. The interface identifier portion of each IPv6 address is not translated; it remains the same on each side of the firewall and visible to anyone who can see the packet header. Additionally, the prefixes are not secure; they can be determined by others.

How NPTv6 Works

When you configure a policy for NPTv6, the Palo Alto Networks firewall performs a static, one-to-one IPv6 translation in both directions. The translation is based on the algorithm described in RFC 6296.

In one use case, the firewall performing NPTv6 is located between an internal network and an external network (such as the Internet) that uses globally routable prefixes. When datagrams are going in the outbound direction, the internal source prefix is replaced with the external prefix; this is known as source translation.

In another use case, when datagrams are going in the inbound direction, the destination prefix is replaced with the internal prefix (known as destination translation). The figure below illustrates destination translation and a characteristic of NPTv6: only the prefix portion of an IPv6 address is translated. The host portion of the address is not translated and remains the same on either side of the firewall. In the figure below, the host identifier is 111::55 on both sides of the firewall.

© 2018 Palo Alto Networks, Inc.



It is important to understand that NPTv6 does not provide security. While you are planning your NPTv6 NAT policies, remember also to configure security policies in each direction.

A NAT or NPTv6 policy rule cannot have both the Source Address and the Translated Address set to Any.

In an environment where you want IPv6 prefix translation, three firewall features work together: NPTv6 NAT policies, security policies, and NDP Proxy.

The firewall does not translate the following:

- Addresses that the firewall has in its Neighbor Discovery (ND) cache.
- The subnet 0xFFFF (in accordance with RFC 6296, Appendix B).
- IP multicast addresses.
- IPv6 addresses with a prefix length of /31 or shorter.
- Link-local addresses. If the firewall is operating in virtual wire mode, there are no IP addresses to translate, and the firewall does not translate link-local addresses.
- Addresses for TCP sessions that authenticate peers using the TCP Authentication Option (RFC 5925).

When using NPTv6, performance for fast path traffic is impacted because NPTv6 is performed in the slow path.

NPTv6 will work with IPSec IPv6 only if the firewall is originating and terminating the tunnel. Transit IPSec traffic would fail because the source and/or destination IPv6 address would be modified. A NAT traversal technique that encapsulates the packet would allow IPSec IPv6 to work with NPTv6.

- Checksum-Neutral Mapping
- Bi-Directional Translation
- NPTv6 Applied to a Specific Service

Checksum-Neutral Mapping

The NPTv6 mapping translations that the firewall performs are checksum-neutral, meaning that "... they result in IP headers that will generate the same IPv6 pseudo-header checksum when the checksum is calculated using the standard Internet checksum algorithm [RFC 1071]." See RFC 6296, Section 2.6, for more information about checksum-neutral mapping.

If you are using NPTv6 to perform destination NAT, you can provide the internal IPv6 address and the external prefix/prefix length of the firewall interface in the syntax of the **`test nptv6`** CLI command. The CLI responds with the checksum-neutral, public IPv6 address to use in your NPTv6 configuration to reach that destination.

Bi-Directional Translation

When you Create an NPTv6 Policy on page 924, the **Bi-directional** option in the **Translated Packet** tab provides a convenient way for you to have the firewall create a corresponding NAT or NPTv6 translation in the opposite direction of the translation you configured. By default, **Bi-directional** translation is disabled.

If you enable Bi-directional translation, it is very important to make sure you have security policies in place to control the traffic in both directions. Without such policies, the Bi-

© 2018 Palo Alto Networks, Inc.

directional feature will allow packets to be automatically translated in both directions, which you might not want.

NPTv6 Applied to a Specific Service

The Palo Alto Networks implementation of NPTv6 offers the ability to filter packets to limit which packets are subject to translation. Keep in mind that NPTv6 does not perform port translation. There is no concept of Dynamic IP and Port (DIPP) translation because NPTv6 translates IPv6 prefixes only. However, you can specify that only packets for a certain service port undergo NPTv6 translation. To do so, Create an NPTv6 Policy on page 924 that specifies a **Service** in the Original Packet.

NDP Proxy

Neighbor Discovery Protocol (NDP) for IPv6 performs functions similar to those provided by Address Resolution Protocol (ARP) for IPv4. RFC 4861 defines Neighbor Discovery for IP version 6 (IPv6). Hosts, routers, and firewalls use NDP to determine the link-layer addresses of neighbors on connected links, to keep track of which neighbors are reachable, and to update neighbors' link-layer addresses that have changed. Peers advertise their own MAC address and IPv6 address, and they also solicit addresses from peers.

NDP also supports the concept of *proxy*, when a node has a neighboring device that is able to forward packets on behalf of the node. The device (firewall) performs the role of NDP Proxy.

Palo Alto Networks firewalls support NDP and NDP Proxy on their interfaces. When you configure the firewall to act as an NDP Proxy for addresses, it allows the firewall to send Neighbor Discovery (ND) advertisements and respond to ND solicitations from peers that are asking for MAC addresses of IPv6 prefixes assigned to devices behind the firewall. You can also configure addresses for which the firewall will not respond to proxy requests (negated addresses).

In fact, NDP is enabled by default, and you need to configure NDP Proxy when you configure NPTv6, for the following reasons:

- The stateless nature of NPTv6 requires a way to instruct the firewall to respond to ND packets sent to specified NDP Proxy addresses, and to not respond to negated NDP Proxy addresses.

  It is recommended that you negate your neighbors' addresses in the NDP Proxy configuration, because NDP Proxy indicates the firewall will reach those addresses behind the firewall, but the neighbors are not behind the firewall.
- NDP causes the firewall to save the MAC addresses and IPv6 addresses of neighbors in its ND cache. (Refer to the figure in NPTv6 and NDP Proxy Example on page 923.) The firewall does not perform NPTv6 translation for addresses that it finds in its ND cache because doing so could introduce a conflict. If the host portion of an address in the cache happens to overlap with the host portion of a neighbor's address, and the prefix in the cache is translated to the same prefix as that of the neighbor (because the egress interface on the firewall belongs to the same subnet as the neighbor), then you would have a translated address that is exactly the same as the legitimate IPv6 address of the neighbor, and a conflict occurs. (If an attempt to perform NPTv6 translation occurs on an address in the ND cache, an informational syslog message logs the event: `NPTv6 Translation Failed.`)

When an interface with NDP Proxy enabled receives an ND solicitation requesting a MAC address for an IPv6 address, the following sequence occurs:

- The firewall searches the ND cache to ensure the IPv6 address from the solicitation is not there. If the address is there, the firewall ignores the ND solicitation.
- If the source IPv6 address is 0, that means the packet is a Duplicate Address Detection packet, and the firewall ignores the ND solicitation.

© 2018 Palo Alto Networks, Inc.

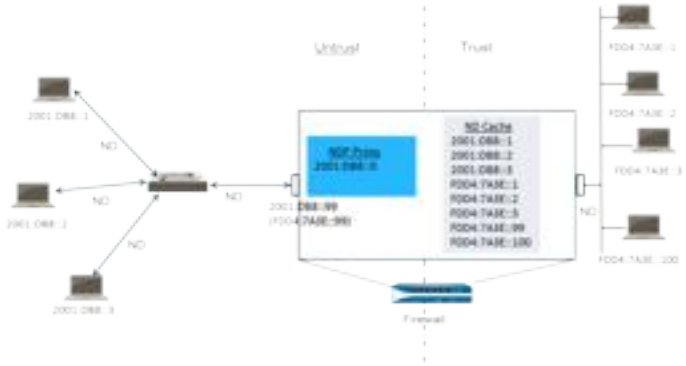
- The firewall does a Longest Prefix Match search of the NDP Proxy addresses and finds the best match to the address in the solicitation. If the Negate field for the match is checked (in the NDP Proxy list), the firewall drops the ND solicitation.
- Only if the Longest Prefix Match search matches, and that matched address is not negated, will the NDP Proxy respond to the ND solicitation. The firewall responds with an ND packet, providing its own MAC address as the MAC address of the next hop toward the queried destination.

In order to successfully support NDP, the firewall does not perform NDP Proxy for the following:

- Duplicate Address Detection (DAD).
- Addresses in the ND cache (because such addresses do not belong to the firewall; they belong to discovered neighbors).

NPTv6 and NDP Proxy Example

The following figure illustrates how NPTv6 and NDP Proxy function together.



- The ND Cache in NPTv6 Example
- The NDP Proxy in NPTv6 Example
- The NPTv6 Translation in NPTv6 Example
- Neighbors in the ND Cache are Not Translated

The ND Cache in NPTv6 Example

In the above example, multiple peers connect to the firewall though a switch, with ND occurring between the peers and the switch, between the switch and the firewall, and between the firewall and the devices on the trust side.

As the firewall learns of peers, it saves their addresses to its ND cache. Trusted peers FDDA:7A3E::1, FDDA:7A3E::2, and FDDA:7A3E::3 are connected to the firewall on the trust side. FDDA:7A3E::99 is the untranslated address of the firewall itself; its public-facing address is 2001:DB8::99. The addresses of the peers on the untrust side have been discovered and appear in the ND cache: 2001:DB8::1, 2001:DB8::2, and 2001:DB8::3.

The NDP Proxy in NPTv6 Example

In our scenario, we want the firewall to act as NDP Proxy for the prefixes on devices behind the firewall. When the firewall is NDP Proxy for a specified set of addresses/ranges/prefixes, and it sees an address from this range in an ND solicitation or advertisement, the firewall will respond as long as a device with that specific address doesn't respond first, the address is not negated in the NDP proxy configuration, and the address is not in the ND cache. The firewall does the prefix translation (described below) and sends the packet to the trust side, where that address might or might not be assigned to a device.

© 2018 Palo Alto Networks, Inc.

In this example, the ND Proxy table contains the network address 2001:DB8::0. When the interface sees an ND for 2001:DB8::100, no other devices on the L2 switch claim the packet, so the proxy range causes the firewall to claim it, and after translation to FDD4:7A3E::100, the firewall sends it out to the trust side.

The NPTv6 Translation in NPTv6 Example

In this example, the **Original Packet** is configured with a **Source Address** of FDD4:7A3E::0 and a **Destination** of **Any**. The **Translated Packet** is configured with the **Translated Address** of 2001:DB8::0.

Therefore, outgoing packets with a source of FDD4:7A3E::0 are translated to 2001:DB8::0. Incoming packets with a destination prefix in the network 2001:DB8::0 are translated to FDD4:7A3E::0.

Neighbors in the ND Cache are Not Translated

In our example, there are hosts behind the firewall with host identifiers :1, :2, and :3. If the prefixes of those hosts are translated to a prefix that exists beyond the firewall, and if those devices also have host identifiers :1, :2, and :3, because the host identifier portion of the address remains unchanged, the resulting translated address would belong to the existing device, and an addressing conflict would result. In order to avoid a conflict with overlapping host identifiers, NPTv6 does not translate addresses that it finds it its ND cache.

Create an NPTv6 Policy

Perform this task when you want to configure a NAT NPTv6 policy to translate one IPv6 prefix to another IPv6 prefix. The prerequisites for this task are:

- Enable IPv6. Select **Device** > **Setup** > **Session**. Click **Edit** and select **IPv6 Firewalling**.
- Configure a Layer 3 Ethernet interface with a valid IPv6 address and with IPv6 enabled. Select **Network** > **Interfaces** > **Ethernet**, select an interface, and on the **IPv6** tab, select **Enable IPv6 on the interface**.
- Create network security policies, because NPTv6 does not provide security.
- Decide whether you want source translation, destination translation, or both.
- Identify the zones to which you want to apply the NPTv6 policy.
- Identify your original and translated IPv6 prefixes.

STEP 1 | Create a new NPTv6 policy.

1. Select **Policies** > **NAT** and click **Add**.
2. On the **General** tab, enter a descriptive **Name** for the NPTv6 policy rule.
3. (Optional) Enter a **Description** and **Tag**.
4. For **NAT Type**, select **NPTv6**.

STEP 2 | Specify the match criteria for incoming packets; packets that match all of the criteria are subject to the NPTv6 translation.

Zones are required for both types of translation.

1. On the **Original Packet** tab, for **Source Zone**, leave **Any** or **Add** the source zone to which the policy applies.
2. Enter the **Destination Zone** to which the policy applies.
3. (Optional) Select a **Destination Interface**.
4. (Optional) Select a **Service** to restrict what type of packets are translated.
5. If you are doing source translation, enter a **Source Address** or select **Any**. The address could be an address object. The following constraints apply to **Source Address** and **Destination Address**:

© 2018 Palo Alto Networks, Inc.

- Prefixes of **Source Address** and **Destination Address** for the **Original Packet** and **Translated Packet** must be in the format xxxx:xxxx::/yy, although leading zeros in the prefix can be dropped.
- The IPv6 address cannot have an interface identifier (host) portion defined.
- The range of supported prefix lengths is /32 to /64.
- The **Source Address** and **Destination Address** cannot both be set to **Any**.

6. If you are doing source translation, you can optionally enter a **Destination Address**. If you are doing destination translation, the **Destination Address** is required. The destination address (an address object is allowed) must be a netmask, not just an IPv6 address and not a range. The prefix length must be a value from /32 to /64, inclusive. For example, 2001:db8::/32.

STEP 3 | Specify the translated packet.

1. On the **Translated Packet** tab, if you want to do source translation, in the Source Address Translation section, for **Translation Type**, select **Static IP**. If you do not want to do source translation, select **None**.
2. If you chose **Static IP**, the **Translated Address** field appears. Enter the translated IPv6 prefix or address object. See the constraints listed in the prior step.

   It is a best practice to configure your Translated Address to be the prefix of the untrust interface address of your firewall. For example, if your untrust interface has the address 2001:1a:1b:1::99/64, make your Translated Address 2001:1a:1b:1::0/64.
3. (Optional) Select **Bi-directional** if you want the firewall to create a corresponding NPTv6 translation in the opposite direction of the translation you configure.

   If you enable Bi-directional translation, it is very important to make sure you have Security policy rules in place to control the traffic in both directions. Without such policy rules, Bi-directional translation allows packets to be automatically translated in both directions, which you might not want.
4. If you want to do destination translation, select **Destination Address Translation**. In the **Translated Address** field, choose an address object from the drop-down or enter your internal destination address.
5. Click **OK**.

STEP 4 | Configure NDP Proxy.

When you configure the firewall to act as an NDP Proxy for addresses, it allows the firewall to send Neighbor Discovery (ND) advertisements and respond to ND solicitations from peers that are asking for MAC addresses of IPv6 prefixes assigned to devices behind the firewall.

1. Select **Network** > **Interfaces** > **Ethernet** and select an interface.
2. On the **Advanced** > **NDP Proxy** tab, select **Enable NDP Proxy** and click **Add**.
3. Enter the **IP Address(es)** for which NDP Proxy is enabled. It can be an address, a range of addresses, or a prefix and prefix length. The order of IP addresses does not matter. These addresses are ideally the same as the Translated Addresses that you configured in an NPTv6 policy.

   If the address is a subnet, the NDP Proxy will respond to all addresses in the subnet, so you should list the neighbors in that subnet with Negate selected, as described in the next step.
4. (Optional) Enter one or more addresses for which you do not want NDP Proxy enabled, and select **Negate**. For example, from an IP address range or prefix range configured in the prior step, you could negate a smaller subset of addresses. It is recommended that you negate the addresses of the neighbors of the firewall.

STEP 5 | Commit the configuration.

© 2018 Palo Alto Networks, Inc.

Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

NAT64

NAT64 provides a way to transition to IPv6 while you still need to communicate with IPv4 networks. When you need to communicate from an IPv6-only network to an IPv4 network, you use NAT64 to translate source and destination addresses from IPv6 to IPv4 and vice versa. NAT64 allows IPv6 clients to access IPv4 servers and allows IPv4 clients to access IPv6 servers. You should understand NAT before configuring NAT64.

- NAT64 Overview
- IPv4-Embedded IPv6 Address
- DNS64 Server
- Path MTU Discovery
- IPv6-Initiated Communication
- Configure NAT64 for IPv6-Initiated Communication
- Configure NAT64 for IPv4-Initiated Communication
- Configure NAT64 for IPv4-Initiated Communication with Port Translation

NAT64 Overview

You can configure two types of NAT64 translation on a Palo Alto Networks firewall; each one is doing a bidirectional translation between the two IP address families:

- The firewall supports stateful NAT64 for IPv6-Initiated Communication on page 929, which maps multiple IPv6 addresses to one IPv4 address, thus preserving IPv4 addresses. (It does not support stateless NAT64, which maps one IPv6 address to one IPv4 address and therefore does not preserve IPv4 addresses.) Configure NAT64 for IPv6-Initiated Communication on page 930.
- The firewall supports IPv4-initiated communication with a static binding that maps an IPv4 address and port number to an IPv6 address. Configure NAT64 for IPv4-Initiated Communication on page 932. It also supports port rewrite, which preserves even more IPv4 addresses by translating an IPv4 address and port number to an IPv6 address with multiple port numbers. Configure NAT64 for IPv4-Initiated Communication with Port Translation on page 934.

A single IPv4 address can be used for NAT44 and NAT64; you don't reserve a pool of IPv4 addresses for NAT64 only.

NAT64 operates on Layer 3 interfaces, subinterfaces, and tunnel interfaces. To use NAT64 on a Palo Alto Networks firewall for IPv6-initiated communication, you must have a third-party DNS64 Server on page 928 or a solution in place to separate the DNS query function from the NAT function. The DNS64 server translates between your IPv6 host and an IPv4 DNS server by encoding the IPv4 address it receives from a public DNS server into an IPv6 address for the IPv6 host.

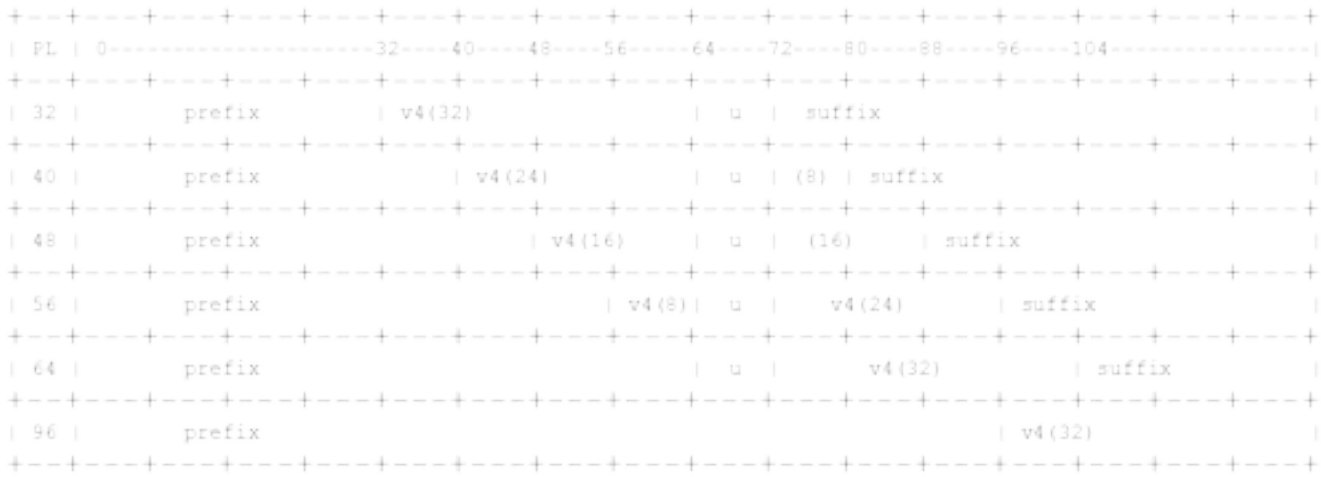
Palo Alto Networks supports the following NAT64 features:

- Hairpinning (NAT U-Turn); additionally, NAT64 prevents hairpinning loop attacks by dropping all incoming IPv6 packets that have a source prefix of 64::/n.
- Translation of TCP/UDP/ICMP packets per RFC 6146 and the firewall makes a best effort to translate other protocols that don't use an application-level gateway (ALG). For example, the firewall can translate a GRE packet. This translation has the same limitation as NAT44: if you don't have an ALG for a protocol that can use a separate control and data channel, the firewall might not understand the return traffic flow.
- Translation between IPv4 and IPv6 of the ICMP length attribute of the original datagram field, per RFC 4884.

© 2018 Palo Alto Networks, Inc.

IPv4-Embedded IPv6 Address

NAT64 uses an IPv4-embedded IPv6 address as described in RFC 6052, IPv6 Addressing of IPv4/IPv6 Translators. An IPv4-embedded IPv6 address is an IPv6 address in which 32 bits have an IPv4 address encoded in them. The IPv6 prefix length (PL in the figure) determines where in the IPv6 address the IPv4 address is encoded, as follows:



The firewall supports translation for /32, /40, /48, /56, /64, and /96 subnets using these prefixes. A single firewall supports multiple prefixes; each NAT64 rule uses one prefix. The prefix can be the Well-Known Prefix (64:FF9B::/96) or a Network-Specific Prefix (NSP) that is unique to the organization that controls the address translator (the DNS64 device). An NSP is usually a network within the organization's IPv6 prefix. The DNS64 device typically sets the u field and suffix to zeros; the firewall ignores those fields.

DNS64 Server

If you need to use a DNS and you want to perform NAT64 translation using IPv6-Initiated Communication, you must use a third-party DNS64 server or other DNS64 solution that is set up with the Well-Known Prefix or your NSP. When an IPv6 host attempts to access an IPv4 host or domain on the internet, the DNS64 server queries an authoritative DNS server for the IPv4 address mapped to that host name. The DNS server returns an Address record (A record) to the DNS64 server containing the IPv4 address for the host name.

The DNS64 server in turn converts the IPv4 address to hexadecimal and encodes it into the appropriate octets of the IPv6 prefix it is set up to use (the Well-Known Prefix or your NSP) based on the prefix length, which results in an IPv4-Embedded IPv6 Address. The DNS64 server sends an AAAA record to the IPv6 host that maps the IPv4-embedded IPv6 address to the IPv4 host name.

Path MTU Discovery

IPv6 does not fragment packets, so the firewall uses two methods to reduce the need to fragment packets:

- When the firewall is translating IPv4 packets in which the DF (don't fragment) bit is zero, that indicates the sender expects the firewall to fragment packets that are too large, but the firewall doesn't fragment packets for the IPv6 network (after translation) because IPv6 doesn't fragment packets. Instead, you can configure the minimum size into which the firewall will fragment IPv4 packets before translating them. The **NAT64 IPv6 Minimum Network MTU** value is this setting, which complies with RFC 6145, IP/ICMP Translation Algorithm. You can set the **NAT64 IPv6 Minimum Network MTU** to its maximum value (**Device** > **Setup** > **Session**), which causes the firewall to fragment IPv4 packets to the IPv6 minimum

© 2018 Palo Alto Networks, Inc.

size before translating them to IPv6. (The **NAT64 IPv6 Minimum Network MTU** does not change the interface MTU.)

- The other method the firewall uses to reduce fragmentation is Path MTU Discovery (PMTUD). In an IPv4-initiated communication, if an IPv4 packet to be translated has the DF bit set and the MTU for the egress interface is smaller than the packet, the firewall uses PMTUD to drop the packet and return an ICMP 'Destination Unreachable - fragmentation needed' message to the source. The source lowers the path MTU for that destination and resends the packet until successive reductions in the path MTU allow packet delivery.

IPv6-Initiated Communication

IPv6-initiated communication to the firewall is similar to source NAT for an IPv4 topology. Configure NAT64 for IPv6-Initiated Communication when your IPv6 host needs to communicate with an IPv4 server.

In the NAT64 policy rule, configure the original source to be an IPv6 host address or Any. Configure the destination IPv6 address as either the Well-Known Prefix or the NSP that the DNS64 server uses. (You do not configure the full IPv6 destination address in the rule.)

If you need to use a DNS, you need to use a DNS64 Server to convert an IPv4 DNS "A" result into an "AAAA" result merged with the NAT64 prefix. If you don't use a DNS, you need to create the address using the IPv4 destination address and the NAT64 prefix configured on the firewall, following RFC 6052 rules.

For environments that use a DNS, the example topology below illustrates communication with the DNS64 Server. The DNS64 server must be set up to use the Well-Known Prefix 64:FF9B::/96 or your Network-Specific Prefix, which must comply with RFC 6052 (/32, /40,/48,/56,/64, or /96).

On the translated side of the firewall, the translation type must be Dynamic IP and Port in order to implement stateful NAT64. You configure the source translated address to be the IPv4 address of the egress interface on the firewall. You do not configure the destination translation field; the firewall translates the address by first finding the prefix length in the original destination address of the rule and then based on the prefix, extracting the encoded IPv4 address from the original destination IPv6 address in the incoming packet.

Before the firewall looks at the NAT64 rule, the firewall must do a route lookup to find the destination security zone for an incoming packet. You must ensure that the NAT64 prefix can be reached through the destination zone assignment because the NAT64 prefix should not be routable by the firewall. The firewall would likely assign the NAT64 prefix to the default route or drop the NAT64 prefix because there is no route for it. The firewall will not find a destination zone because the NAT64 prefix is not in its routing table, associated with an egress interface and zone.

You must also configure a tunnel interface (with no termination point). You apply the NAT64 prefix to the tunnel and apply the appropriate zone to ensure that IPv6 traffic with the NAT64 prefix is assigned to the proper destination zone. The tunnel also has the advantage of dropping IPv6 traffic with the NAT64 prefix if the traffic does not match the NAT64 rule.Your configured routing protocol on the firewall looks up the IPv6 prefix in its routing table to find the destination zone and then looks at the NAT64 rule.

The following figure illustrates the role of the DNS64 server in the name resolution process. In this example, the DNS64 server is configured to use Well-Known Prefix 64:FF9B::/96.

1. A user at the IPv6 host enters the URL www.abc.com, which generates a name server lookup (nslookup) to the DNS64 server.

2. The DNS64 Server sends an nslookup to the public DNS server for www.abc.com, requesting its IPv4 address.

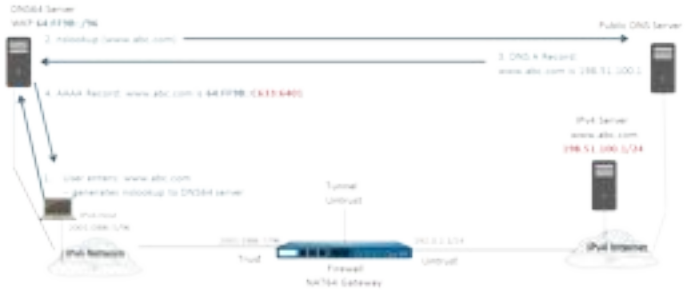
3. The DNS server returns an A record that provides the IPv4 address to the DNS64 server.

4. The DNS64 server sends an AAAA record to the IPv6 user, converting the IPv4 dotted decimal address 198.51.100.1 into C633:6401 hexadecimal and embedding it into its own IPv6 prefix, 64:FF9B::/96.

© 2018 Palo Alto Networks, Inc.

[198 = C6 hex; 51 = 33 hex; 100 = 64 hex; 1 = 01 hex.] The result is IPv4-Embedded IPv6 Address 64:FF9B::C633:6401.

Keep in mind that in a /96 prefix, the IPv4 address is the last four octets encoded in the IPv6 address. If the DNS64 server uses a /32, /40, /48, /56 or /64 prefix, the IPv4 address is encoded as shown in RFC 6052.



Upon the transparent name resolution, the IPv6 host sends a packet to the firewall containing its IPv6 source address and destination IPv6 address 64:FF9B::C633:6401 as determined by the DNS64 server. The firewall performs the NAT64 translation based on your NAT64 rule.



Configure NAT64 for IPv6-Initiated Communication

This configuration task and its addresses correspond to the figures in IPv6-Initiated Communication.

STEP 1 | Enable IPv6 to operate on the firewall.

1. Select **Device** > **Setup** > **Session** and edit the Session Settings.
2. Select **Enable IPv6 Firewalling**.
3. Click **OK**.

STEP 2 | Create an address object for the IPv6 destination address (pre-translation).

1. Select **Objects** > **Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64-IPv4 Server.
3. For **Type**, select **IP Netmask** and enter the IPv6 prefix with a netmask that is compliant with RFC 6052 (/32, /40, /48, /56, /64, or /96). This is either the Well-Known Prefix or your Network-Specific Prefix that is configured on the DNS64 Server.

   For this example, enter 64:FF9B::/96.

   The source and destination must have the same netmask (prefix length).

   (You don't enter a full destination address because, based on the prefix length, the firewall extracts the encoded IPv4 address from the original destination IPv6 address in the incoming packet. In this example, the prefix in the incoming packet is encoded with C633:6401 in hexadecimal, which is the IPv4 destination address 198.51.100.1.)

© 2018 Palo Alto Networks, Inc.

4. Click **OK**.

STEP 3 | (Optional) Create an address object for the IPv6 source address (pre-translation).

1. Select **Objects** > **Addresses** and click **Add**.
2. Enter a **Name** for the object.
3. For **Type**, select **IP Netmask** and enter the address of the IPv6 host, in this example, 2001:DB8::5/96.
4. Click **OK**.

STEP 4 | (Optional) Create an address object for the IPv4 source address (translated).

1. Select **Objects** > **Addresses** and click **Add**.
2. Enter a **Name** for the object.
3. For **Type**, select **IP Netmask** and enter the IPv4 address of the firewall's egress interface, in this example, 192.0.2.1.
4. Click **OK**.

STEP 5 | Create the NAT64 rule.

1. Select **Policies** > **NAT** and click **Add**.
2. On the **General** tab, enter a **Name** for the NAT64 rule, for example, nat64_ipv6_init.
3. (Optional) Enter a **Description**.
4. For **NAT Type**, select **nat64**.

STEP 6 | Specify the original source and destination information.

1. For the **Original Packet**, **Add** the **Source Zone**, likely a trusted zone.
2. Select the **Destination Zone**, in this example, the Untrust zone.
3. (Optional) Select a **Destination Interface** or the default (**any**).
4. For **Source Address**, select **Any**or **Add** the address object you created for the IPv6 host.
5. For **Destination Address**, **Add** the address object you created for the IPv6 destination address, in this example, nat64-IPv4 Server.
6. (Optional) For **Service**, select **any**.

STEP 7 | Specify the translated packet information.

1. For the **Translated Packet**, in **Source Address Translation**, for **Translation Type**, select **Dynamic IP and Port**.
2. For **Address Type**, do one of the following:
   - Select **Translated Address** and **Add** the address object you created for the IPv4 source address.
   - Select **Interface Address**, in which case the translated source address is the IP address and netmask of the firewall's egress interface. For this choice, select an **Interface** and optionally an **IP Address** if the interface has more than one IP address.
3. Leave **Destination Address Translation** unselected. (The firewall extracts the IPv4 address from the IPv6 prefix in the incoming packet, based on the prefix length specified in the original destination of the NAT64 rule.)
4. Click **OK** to save the NAT64 policy rule.

STEP 8 | Configure a tunnel interface to emulate a loopback interface with a netmask other than 128.

1. Select **Network** > **IPSec Tunnels** and **Add** a tunnel.
2. On the **General** tab, enter a **Name** for the tunnel.
3. For the **Tunnel Interface**, select **New Tunnel Interface**.
4. In the **Interface Name** field, enter a numeric suffix, such as .2.
5. On the **Config** tab, select the **Virtual Router** where you are configuring NAT64.

© 2018 Palo Alto Networks, Inc.

6. For **Security Zone**, select the destination zone associated with the IPv4 server destination (Trust zone).
7. On the **IPv6** tab, select **Enable IPv6 on the interface**.
8. Click **Add** and for the **Address**, select **New Address**.
9. Enter a **Name** for the address.
10. (Optional) Enter a **Description** for the tunnel address.
11. For **Type**, select **IP Netmask** and enter your IPv6 prefix and prefix length, in this example, 64:FF9B::/96.
12. Click **OK** to save the new address.
13. Select **Enable address on interface** and click **OK**.
14. Click **OK** to save the tunnel interface.
15. Click **OK** to save the tunnel.

STEP 9 | Create a security policy to allow NAT traffic from the trust zone.

1. Select **Policies** > **Security** and **Add** a rule **Name**.
2. Select **Source** and **Add** a **Source Zone**; select **Trust**.
3. For **Source Address**, select **Any**.
4. Select **Destination** and **Add** a **Destination Zone**; select **Untrust**.
5. For **Application**, select **Any**.
6. For **Actions**, select **Allow**.
7. Click **OK**.

STEP 10 | Commit your changes.

Click **Commit**.

STEP 11 | Troubleshoot or view a NAT64 session.

```
> show session id <session-id>
```

Configure NAT64 for IPv4-Initiated Communication

IPv4-initiated communication to an IPv6 server is similar to destination NAT in an IPv4 topology. The destination IPv4 address maps to the destination IPv6 address through a one-to-one, static IP translation (not a many-to-one translation).

The firewall encodes the source IPv4 address into Well-Known Prefix 64:FF9B::/96 as defined in RFC 6052. The translated destination address is the actual IPv6 address. The use case for IPv4-initiated communication is typically when an organization is providing access from the public, untrust zone to an IPv6 server in the organization's DMZ zone. This topology does not use a DNS64 server.



STEP 1 | Enable IPv6 to operate on the firewall.

1. Select **Device** > **Setup** > **Session** and edit the Session Settings.
2. Select **Enable IPv6 Firewalling**.

© 2018 Palo Alto Networks, Inc.

3. Click **OK**.

STEP 2 | (Optional) When an IPv4 packet has its DF bit set to zero (and because IPv6 does not fragment packets), ensure the translated IPv6 packet does not exceed the path MTU for the destination IPv6 network.

1. Select **Device** > **Setup** > **Session** and edit Session Settings.
2. For **NAT64 IPv6 Minimum Network MTU**, enter the smallest number of bytes into which the firewall will fragment IPv4 packets for translation to IPv6 (range is 1280-9216, default is 1280).



If you don't want the firewall to fragment an IPv4 packet prior to translation, set the MTU to 9216. If the translated IPv6 packet still exceeds this value, the firewall drops the packet and issues an ICMP packet indicating destination unreachable - fragmentation needed.

3. Click **OK**.

STEP 3 | Create an address object for the IPv4 destination address (pre-translation).

1. Select **Objects** > **Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_ip4server.
3. For **Type**, select **IP Netmask** and enter the IPv4 address and netmask of the firewall interface in the Untrust zone. This example uses 198.51.19.1/24.
4. Click **OK**.

STEP 4 | Create an address object for the IPv6 source address (translated).

1. Select **Objects** > **Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_ip6source.
3. For **Type**, select **IP Netmask** and enter the NAT64 IPv6 address with a netmask that is compliant with RFC 6052 (/32, /40, /48, /56, /64, or /96).

   For this example, enter 64:FF9B::/96.

   (The firewall encodes the prefix with the IPv4 source address 192.1.2.8, which is C001:0208 in hexadecimal.)
4. Click **OK**.

STEP 5 | Create an address object for the IPv6 destination address (translated).

1. Select **Objects** > **Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_server_2.
3. For **Type**, select **IP Netmask** and enter the IPv6 address of the IPv6 server (destination). This example uses 2001:DB8::2/64.



The source and destination must have the same netmask (prefix length).

4. Click **OK**.

STEP 6 | Create the NAT64 rule.

1. Select **Policies** > **NAT** and click **Add**.
2. On the **General** tab, enter a **Name** for the NAT64 rule, for example, nat64_ipv4_init.
3. For **NAT Type**, select **nat64**.

STEP 7 | Specify the original source and destination information.

1. For the **Original Packet**, **Add** the **Source Zone**, likely an untrust zone.
2. Select the **Destination Zone**, likely a trust or DMZ zone.

© 2018 Palo Alto Networks, Inc.

3. For **Source Address**, select **Any**or **Add** the address object for the IPv4 host.
4. For **Destination Address**, **Add** the address object for the IPv4 destination, in this example, nat64_ip4server.
5. For **Service**, select **any**.

STEP 8 | Specify the translated packet information.

1. For the **Translated Packet**, in the **Source Address Translation**, **Translation Type**, select **Static IP**.
2. For **Translated Address**, select the source translated address object you created, nat64_ip6source.
3. For **Destination Address Translation**, for **Translated Address**, specify a single IPv6 address (the address object, in this example, nat64_server_2, or the IPv6 address of the server).
4. Click **OK**.

STEP 9 | Create a security policy to allow the NAT traffic from the Untrust zone.

1. Select **Policies** > **Security** and **Add** a rule **Name**.
2. Select **Source** and **Add** a **Source Zone**; select **Untrust**.
3. For **Source Address**, select **Any**.
4. Select **Destination** and **Add** a **Destination Zone**; select **DMZ**.
5. For **Actions**, select **Allow**.
6. Click **OK**.

STEP 10 | Commit your changes.

Click **Commit**.

STEP 11 | Troubleshoot or view a NAT64 session.

```
> show session id <session-id>
```

Configure NAT64 for IPv4-Initiated Communication with Port Translation

This task builds on the task to Configure NAT64 for IPv4-Initiated Communication, but the organization controlling the IPv6 network prefers to translate the public destination port number to an internal destination port number and thereby keep it private from users on the IPv4 untrust side of the firewall. In this example, port 8080 is translated to port 80. To do that, in the Original Packet of the NAT64 policy rule, create a new Service that specifies the destination port is 8080. For the Translated Packet, the translated port is 80.



STEP 1 | Enable IPv6 to operate on the firewall.

1. Select **Device** > **Setup** > **Session** and edit the Session Settings.
2. Select **Enable IPv6 Firewalling**.
3. Click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 2 | (Optional) When an IPv4 packet has its DF bit set to zero (and because IPv6 does not fragment packets), ensure the translated IPv6 packet does not exceed the path MTU for the destination IPv6 network.

1. Select **Device** > **Setup** > **Session** and edit Session Settings.
2. For **NAT64 IPv6 Minimum Network MTU**, enter the smallest number of byes into which the firewall will fragment IPv4 packets for translation to IPv6 (range is 1280-9216, default is 1280).

   If you don't want the firewall to fragment an IPv4 packet prior to translation, set the MTU to 9216. If the translated IPv6 packet still exceeds this value, the firewall drops the packet and issues an ICMP packet indicating destination unreachable - fragmentation needed.
3. Click **OK**.

STEP 3 | Create an address object for the IPv4 destination address (pre-translation).

1. Select **Objects** > **Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_ip4server.
3. For **Type**, select **IP Netmask** and enter the IPv4 address and netmask of the firewall interface in the Untrust zone. This example uses 198.51.19.1/24.
4. Click **OK**.

STEP 4 | Create an address object for the IPv6 source address (translated).

1. Select **Objects** > **Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_ip6source.
3. For **Type**, select **IP Netmask** and enter the NAT64 IPv6 address with a netmask that is compliant with RFC 6052 (/32, /40, /48, /56, /64, or /96).

   For this example, enter 64:FF9B::/96.

   (The firewall encodes the prefix with the IPv4 source address 192.1.2.8, which is C001:0208 in hexadecimal.)
4. Click **OK**.

STEP 5 | Create an address object for the IPv6 destination address (translated).

1. Select **Objects** > **Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_server_2.
3. For **Type**, select **IP Netmask** and enter the IPv6 address of the IPv6 server (destination). This example uses 2001:DB8::2/64.

   The source and destination must have the same netmask (prefix length).
4. Click **OK**.

STEP 6 | Create the NAT64 rule.

1. Select **Policies** > **NAT** and click **Add**.
2. On the **General** tab, enter a **Name** for the NAT64 rule, for example, nat64_ipv4_init.
3. For **NAT Type**, select **nat64**.

STEP 7 | Specify the original source and destination information, and create a service to limit the translation to a single ingress port number.

1. For the **Original Packet**, **Add** the **Source Zone**, likely an untrust zone.
2. Select the **Destination Zone**, likely a trust or DMZ zone.
3. For **Service**, select New **Service**.

© 2018 Palo Alto Networks, Inc.

4. Enter a **Name** for the Service, such as Port_8080.
5. Select **TCP** as the **Protocol**.
6. For **Destination Port**, enter 8080.
7. Click **OK** to save the Service.
8. For **Source Address**, select **Any**or **Add** the address object for the IPv4 host.
9. For **Destination Address**, **Add** the address object for the IPv4 destination, in this example, nat64_ip4server.

STEP 8 | Specify the translated packet information.

1. For the **Translated Packet**, in the **Source Address Translation**, **Translation Type**, select **Static IP**.
2. For **Translated Address**, select the source translated address object you created, nat64_ip6source.
3. For **Destination Address Translation**, for **Translated Address**, specify a single IPv6 address (the address object, in this example, nat64_server_2, or the IPv6 address of the server).
4. Specify the private destination **Translated Port** number to which the firewall translates the public destination port number, in this example, 80.
5. Click **OK**.

STEP 9 | Create a security policy to allow the NAT traffic from the Untrust zone.

1. Select **Policies** > **Security** and **Add** a rule **Name**.
2. Select **Source** and **Add** a **Source Zone**; select **Untrust**.
3. For **Source Address**, select **Any**.
4. Select **Destination** and **Add** a **Destination Zone**; select **DMZ**.
5. For **Actions**, select **Allow**.
6. Click **OK**.

STEP 10 | Commit your changes.

Click **Commit**.

STEP 11 | Troubleshoot or view a NAT64 session.

```
> show session id <session-id>
```


© 2018 Palo Alto Networks, Inc.

ECMP

Equal Cost Multiple Path (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, if there are multiple equal-cost routes to the same destination, the virtual router chooses one of those routes from the routing table and adds it to its forwarding table; it will not use any of the other routes unless there is an outage in the chosen route.

Enabling ECMP functionality on a virtual router allows the firewall to have up to four equal-cost paths to a destination in its forwarding table, allowing the firewall to:

- Load balance flows (sessions) to the same destination over multiple equal-cost links.
- Efficiently use all available bandwidth on links to the same destination rather than leave some links unused.
- Dynamically shift traffic to another ECMP member to the same destination if a link fails, rather than having to wait for the routing protocol or RIB table to elect an alternative path/route. This can help reduce downtime when links fail.

For information about ECMP path selection when an HA peer fails, see ECMP in Active/Active HA Mode.

The following sections describe ECMP and how to configure it.

- ECMP Load-Balancing Algorithms
- ECMP Model, Interface, and IP Routing Support
- Configure ECMP on a Virtual Router
- Enable ECMP for Multiple BGP Autonomous Systems
- Verify ECMP

ECMP Load-Balancing Algorithms

Let's suppose the Routing Information Base (RIB) of the firewall has multiple equal-cost paths to a single destination. The maximum number of equal-cost paths defaults to 2. ECMP chooses the best two equal-cost paths from the RIB to copy to the Forwarding Information Base (FIB). ECMP then determines, based on the load-balancing method, which of the two paths in the FIB that the firewall will use for the destination during this session.

ECMP load balancing is done at the session level, not at the packet level—the start of a new session is when the firewall (ECMP) chooses an equal-cost path. The equal-cost paths to a single destination are considered ECMP path members or ECMP group members. ECMP determines which one of the multiple paths to a destination in the FIB to use for an ECMP flow, based on which load-balancing algorithm you set. A virtual router can use only one load-balancing algorithm.

Enabling, disabling, or changing ECMP on an existing virtual router causes the system to restart the virtual router, which might cause existing sessions to be terminated.

The four algorithm choices emphasize different priorities, as follows:

- **Hash-based algorithms prioritize session stickiness**—The **IP Modulo** and **IP Hash** algorithms use hashes based on information in the packet header, such as source and destination address. Because the header of each flow in a given session contains the same source and destination information, these options prioritize session *stickiness*. If you choose the **IP Hash** algorithm, the hash can be based on the source and destination addresses, or the hash can be based on the source address only (in PAN-OS 8.0.3 and later releases). Using an IP hash based on only the source address causes all sessions belonging to the same source IP address to always take the same path from available multiple paths. Thus the path is considered sticky and is easier to troubleshoot if necessary. You can optionally set a **Hash Seed** value to further randomize load balancing if you have a large number of sessions to the same destination and they're not being distributed evenly over the ECMP links.

© 2018 Palo Alto Networks, Inc.

- **Balanced algorithm prioritizes load balancing**—The **Balanced Round Robin** algorithm distributes incoming sessions equally across the links, favoring load balancing over session stickiness. (Round robin indicates a sequence in which the least recently chosen item is chosen.) In addition, if new routes are added or removed from an ECMP group (for example if a path in the group goes down), the virtual router will re-balance the sessions across links in the group. Additionally, if the flows in a session have to switch routes due to an outage, when the original route associated with the session becomes available again, the flows in the session will revert to the original route when the virtual router once again re-balances the load.
- **Weighted algorithm prioritizes link capacity and/or speed**—As an extension to the ECMP protocol standard, the Palo Alto Networks implementation provides for a **Weighted Round Robin** load-balancing option that takes into account differing link capacities and speeds on the egress interfaces of the firewall. With this option, you can assign **ECMP Weights** (range is 1-255; default is 100) to the interfaces based on link performance using factors such as link capacity, speed, and latency to ensure that loads are balanced to fully leverage the available links.

  For example, suppose the firewall has redundant links to an ISP: ethernet1/1 (100 Mbps) and ethernet1/8 (200 Mbps). Although these are equal-cost paths, the link via ethernet1/8 provides greater bandwidth and therefore can handle a greater load than the ethernet1/1 link. Therefore, to ensure that the load-balancing functionality takes into account link capacity and speed, you might assign ethernet1/8 a weight of 200 and ethernet1/1 a weight of 100. The 2:1 weight ratio causes the virtual router to send twice as many sessions to ethernet1/8 as it sends to ethernet1/1. However, because the ECMP protocol is inherently session-based, when using the **Weighted Round Robin** algorithm, the firewall will be able to load balance across the ECMP links only on a best-effort basis.

  Keep in mind that ECMP weights are assigned to interfaces to determine load balancing (to influence which *equal-cost* path is chosen), not for route selection (a route choice from routes that could have different costs).

  Assign lower-speed or lower-capacity links with a lower weight. Assign higher-speed or higher-capacity links with a higher weight. In this manner, the firewall can distribute sessions based on these ratios, rather than overdrive a low-capacity link that is one of the equal-cost paths.

ECMP Model, Interface, and IP Routing Support

ECMP is supported on all Palo Alto Networks firewall models, with hardware forwarding support on the PA-7000 Series, PA-5000 Series, PA-3060 firewalls, and PA-3050 firewalls. PA-3020 firewalls, PA-500 firewalls, PA-200 firewalls, and VM-Series firewalls support ECMP through software only. Performance is affected for sessions that cannot be hardware offloaded.

ECMP is supported on Layer 3, Layer 3 subinterface, VLAN, tunnel, and Aggregated Ethernet interfaces.

ECMP can be configured for static routes and any of the dynamic routing protocols the firewall supports.

ECMP affects the route table capacity because the capacity is based on the number of paths, so an ECMP route with four paths will consume four entries of route table capacity. ECMP implementation might slightly decrease the route table capacity because more memory is being used by session-based tags to map traffic flows to particular interfaces.

Virtual router-to-virtual router routing using static routes does not support ECMP.

Configure ECMP on a Virtual Router

Use the following procedure to enable ECMP on a virtual router. The prerequisites are to:

- Specify the interfaces that belong to a virtual router (**Network** > **Virtual Routers** > **Router Settings** > **General**).

© 2018 Palo Alto Networks, Inc.

- Specify the IP routing protocol.

Enabling, disabling, or changing ECMP for an existing virtual router causes the system to restart the virtual router, which might cause sessions to be terminated.

STEP 1 | Enable ECMP for a virtual router.

1. Select **Network** > **Virtual Routers** and select the virtual router on which to enable ECMP.
2. Select **Router Settings** > **ECMP** and select **Enable**.

STEP 2 | (Optional) Enable symmetric return of packets from server to client.

Select **Symmetric Return** to cause return packets to egress out the same interface on which the associated ingress packets arrived. That is, the firewall will use the ingress interface on which to send return packets, rather than use the ECMP interface. The **Symmetric Return** setting overrides load balancing. This behavior occurs only for traffic flows from the server to the client.

STEP 3 | Specify the maximum number of equal-cost paths (to a destination network) that can be copied from the Routing Information Base (RIB) to the Forwarding Information Base (FIB).

For **Max Path** allowed, enter **2**, **3**, or **4**. Default: 2.

STEP 4 | Select the load-balancing algorithm for the virtual router. For more information on load-balancing methods and how they differ, see ECMP Load-Balancing Algorithms.

For **Load Balance**, select one of the following options from the **Method** drop-down:

- **IP Modulo** (default)—Uses a hash of the source and destination IP addresses in the packet header to determine which ECMP route to use.
- **IP Hash**—There are two IP hash methods that determine which ECMP route to use (select hash options in Step 5):
  - Use a hash of the source address (available in PAN-OS 8.0.3 and later releases).
  - Use a hash of the source and destination IP addresses (the default IP hash method).
- **Balanced Round Robin**—Uses round robin among the ECMP paths and re-balances paths when the number of paths changes.
- **Weighted Round Robin**—Uses round robin and a relative weight to select from among ECMP paths. Specify the weights in Step 6 below.

STEP 5 | (IP Hash only) Configure IP Hash options.

If you selected **IP Hash** as the **Method**:

1. Select **Use Source Address Only** (available in PAN-OS 8.0.3 and later releases) if you want to ensure all sessions belonging to the same source IP address always take the same path from available multiple paths. This IP hash option provides path stickiness and eases troubleshooting. If you don't select this option or you're using a release prior to PAN-OS 8.0.3, the IP hash is based on the source and destination IP addresses (the default IP hash method).

   If you select Use Source Address Only, you shouldn't push the configuration from Panorama to firewalls running PAN-OS 8.0.2, 8.0.1, or 8.0.0.

2. Select **Use Source/Destination Ports** if you want to use source or destination port numbers in the **IP Hash** calculation.

   Enabling this option along with Use Source Address Only will randomize path selection even for sessions belonging to the same source IP address.

© 2018 Palo Alto Networks, Inc.

3. Enter a **Hash Seed** value (an integer with a maximum of nine digits). Specify a **Hash Seed** value to further randomize load balancing. Specifying a hash seed value is useful if you have a large number of sessions with the same tuple information.

STEP 6 | (Weighted Round Robin only) Define a weight for each interface in the ECMP group.

If you selected **Weighted Round Robin** as the **Method**, define a weight for each of the interfaces that are the egress points for traffic to be routed to the same destinations (that is, interfaces that are part of an ECMP group, such as the interfaces that provide redundant links to your ISP or interfaces to the core business applications on your corporate network).

The higher the weight, the more often that equal-cost path will be selected for a new session.

Give higher speed links a higher weight than a slower links so that more of the ECMP traffic goes over the faster link.

1. Create an ECMP group by clicking **Add** and selecting an **Interface** from the drop-down.
2. **Add** the other interfaces in the ECMP group.
3. Click on **Weight** and specify the relative weight for each interface (range is 1-255; default is 100).

STEP 7 | Save the configuration.

1. Click **OK**.
2. At the ECMP Configuration Change prompt, click **Yes** to restart the virtual router. Restarting the virtual router might cause existing sessions to be terminated.

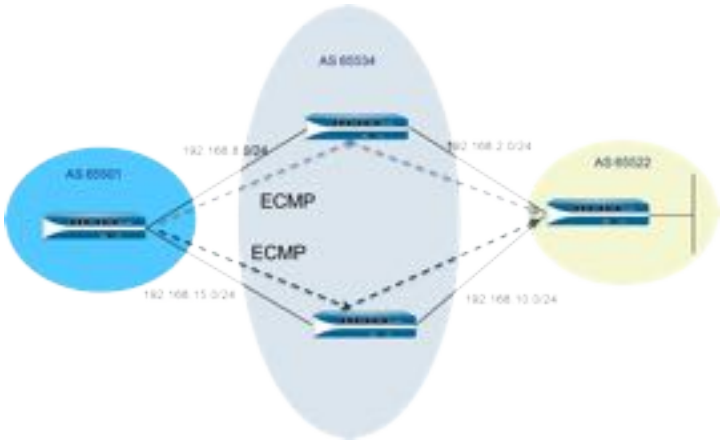
This message displays only if you are modifying an existing virtual router with ECMP.

STEP 8 | Commit your changes.

Commit the configuration.

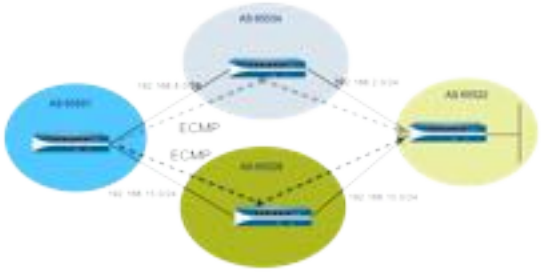
Enable ECMP for Multiple BGP Autonomous Systems

Perform the following task if you have BGP configured, and you want to enable ECMP over multiple autonomous systems. This task presumes that BGP is already configured. In the following figure, two ECMP paths to a destination go through two firewalls belonging to a single ISP in a single BGP autonomous system.



© 2018 Palo Alto Networks, Inc.

In the following figure, two ECMP paths to a destination go through two firewalls belonging to two different ISPs in different BGP autonomous systems.



STEP 1 | Configure ECMP.

See Configure ECMP on a Virtual Router.

STEP 2 | For BGP routing, enable ECMP over multiple autonomous systems.

1. Select **Network** > **Virtual Routers** and select the virtual router on which to enable ECMP for multiple BGP autonomous systems.
2. Select **BGP** > **Advanced** and select **ECMP Multiple AS Support**.

STEP 3 | Commit your changes.

Click **OK** and **Commit**.

Verify ECMP

A virtual router configured for ECMP indicates in the Forwarding Information Base (FIB) table which routes are ECMP routes. An ECMP flag (E) for a route indicates that it is participating in ECMP for the egress interface to the next hop for that route. To verify ECMP, use the following procedure to look at the FIB and confirm that some routes are equal-cost multiple paths.

STEP 1 | Select **Network** > **Virtual Routers**.

STEP 2 | In the row of the virtual router for which you enabled ECMP, click **More Runtime Stats**.

STEP 3 | Select **Routing** > **Forwarding Table** to see the FIB.

In the table, multiple routes to the same Destination (out a different Interface) have the E flag. An asterisk (*) denotes the preferred path for the ECMP group.

© 2018 Palo Alto Networks, Inc.

LLDP

Palo Alto Networks firewalls support Link Layer Discovery Protocol (LLDP), which functions at the link layer to discover neighboring devices and their capabilities. LLDP allows the firewall and other network devices to send and receive LLDP data units (LLDPDUs) to and from neighbors. The receiving device stores the information in a MIB, which the Simple Network Management Protocol (SNMP) can access. LLDP makes troubleshooting easier, especially for virtual wire deployments where the firewall would typically go undetected by a ping or traceroute.

- LLDP Overview
- Supported TLVs in LLDP
- LLDP Syslog Messages and SNMP Traps
- Configure LLDP
- View LLDP Settings and Status
- Clear LLDP Statistics

LLDP Overview

LLDP operates at Layer 2 of the OSI model, using MAC addresses. An LLDPDU is a sequence of type-length-value (TLV) elements encapsulated in an Ethernet frame. The IEEE 802.1AB standard defines three MAC addresses for LLDPDUs: 01-80-C2-00-00-0E, 01-80-C2-00-00-03, and 01-80-C2-00-00-00.

The Palo Alto Networks firewall supports only one MAC address for transmitting and receiving LLDP data units: 01-80-C2-00-00-0E. When transmitting, the firewall uses 01-80-C2-00-00-0E as the destination MAC address. When receiving, the firewall processes datagrams with 01-80-C2-00-00-0E as the destination MAC address. If the firewall receives either of the other two MAC addresses for LLDPDUs on its interfaces, the firewall takes the same forwarding action it took prior to this feature, as follows:

- If the interface type is vwire, the firewall forwards the datagram to the other port.
- If the interface type is L2, the firewall floods the datagram to the rest of the VLAN.
- If the interface type is L3, the firewall drops the datagrams.

Panorama, the GlobalProtect Mobile Security Manager, and the WildFire appliance are not supported.

Interface types that do not support LLDP are TAP, high availability (HA), Decrypt Mirror, virtual wire/vlan/L3 subinterfaces, and PA-7000 Series Log Processing Card (LPC) interfaces.

An LLDP Ethernet frame has the following format:

Preamble	Destination MAC	Source MAC	Ethertype	Chassis ID TLV	Port ID TLV	Time To Live TLV	Optional TLVs	End of LLDPDU TLV	Frame Check Sequence
	01:80:C2:00:00:0E or 01:80:C2:00:00:03 or 01:80:C2:00:00:00	Station's Address	0x88CC	Type=1	Type=2	Type=3	Zero or more complete TLVs	Type=0, Length=0	

Within the LLDP Ethernet frame, the TLV structure has the following format:

© 2018 Palo Alto Networks, Inc.

TLV Type	TLV Information String Length	TLV Information String
7 bits	9 bits	0-511 octets

Supported TLVs in LLDP

LLDPDUs include mandatory and optional TLVs. The following table lists the mandatory TLVs that the firewall supports:

Mandatory TLVs	TLV Type	Description
Chassis ID TLV	1	Identifies the firewall chassis. Each firewall must have exactly one unique Chassis ID. The Chassis ID subtype is 4 (MAC address) on Palo Alto Networks models will use the MAC address of Eth0 to ensure uniqueness.
Port ID TLV	2	Identifies the port from which the LLDPDU is sent. Each firewall uses one Port ID for each LLDPDU message transmitted. The Port ID subtype is 5 (interface name) and uniquely identifies the transmitting port. The firewall uses the interface's ifname as the Port ID.
Time-to-live (TTL) TLV	3	Specifies how long (in seconds) LLDPDU information received from the peer is retained as valid in the local firewall (range is 0-65,535). The value is a multiple of the LLDP Hold Time Multiplier. When the TTL value is 0, the information associated with the device is no longer valid and the firewall removes that entry from the MIB.
End of LLDPDU TLV	0	Indicates the end of the TLVs in the LLDP Ethernet frame.

The following table lists the optional TLVs that the Palo Alto Networks firewall supports:

Optional TLVs	TLV Type	Purpose and Notes Regarding Firewall Implementation
Port Description TLV	4	Describes the port of the firewall in alpha-numeric format. The ifAlias object is used.
System Name TLV	5	Configured name of the firewall in alpha-numeric format. The sysName object is used.
System Description TLV	6	Describes the firewall in alpha-numeric format. The sysDescr object is used.
System Capabilities	7	Describes the deployment mode of the interface, as follows: • An L3 interface is advertised with router (bit 6) capability and the "other" bit (bit 1).

© 2018 Palo Alto Networks, Inc.

Optional TLVs	TLV Type	Purpose and Notes Regarding Firewall Implementation
		• An L2 interface is advertised with MAC Bridge (bit 3) capability and the “other” bit (bit 1). • A virtual wire interface is advertised with Repeater (bit 2) capability and the “other” bit (bit 1).
Management Address	8	One or more IP addresses used for firewall management, as follows: • IP address of the management (MGT) interface • IPv4 and/or IPv6 address of the interface • Loopback address • User-defined address entered in the management address field If no management IP address is provided, the default is the MAC address of the transmitting interface. Included is the interface number of the management address specified. Also included is the OID of the hardware interface with the management address specified (if applicable). If more than one management address is specified, they will be sent in the order they are specified, starting at the top of the list. A maximum of four Management Addresses are supported. This is an optional parameter and can be left disabled.

LLDP Syslog Messages and SNMP Traps

The firewall stores LLDP information in MIBs, which an SNMP Manager can monitor. If you want the firewall to send SNMP trap notifications and syslog messages about LLDP events, you must enable **SNMP Syslog Notification** in an LLDP profile.

Per RFC 5424, The Syslog Protocol, and RFC 1157, A Simple Network Management Protocol, LLDP sends syslog and SNMP trap messages when MIB changes occur. These messages are rate-limited by the **Notification Interval**, an LLDP global setting that defaults to 5 seconds and is configurable.

Because the LLDP syslog and SNMP trap messages are rate-limited, some LLDP information provided to those processes might not match the current LLDP statistics seen when you View the LLDP status information. This is normal, expected behavior.

A maximum of 5 MIBs can be received per interface (Ethernet or AE). Each different source has one MIB. If this limit is exceeded, the error message `tooManyNeighbors` is triggered.

Configure LLDP

To configure LLDP and create an LLDP profile, you must be a superuser or device administrator (deviceadmin). A firewall interface supports a maximum of five LLDP peers.

STEP 1 | Enable LLDP on the firewall.

Select **Network** > **LLDP** and edit the LLDP General section; select **Enable**.

STEP 2 | (Optional) Change LLDP global settings.

1. For **Transmit Interval (sec)**, specify the interval (in seconds) at which LLDPDUs are transmitted. Default: 30 seconds. Range: 1-3600 seconds.

© 2018 Palo Alto Networks, Inc.

2. For **Transmit Delay (sec)**, specify the delay time (in seconds) between LLDP transmissions sent after a change is made in a TLV element. The delay helps to prevent flooding the segment with LLDPDUs if many network changes spike the number of LLDP changes, or if the interface flaps. The **Transmit Delay** must be less than the **Transmit Interval**. Default: 2 seconds. Range: 1-600 seconds.
3. For **Hold Time Multiple**, specify a value that is multiplied by the **Transmit Interval** to determine the total TTL Hold Time. Default: 4. Range: 1-100. The maximum TTL Hold Time is 65535 seconds, regardless of the multiplier value.
4. For **Notification Interval**, specify the interval (in seconds) at which LLDP Syslog Messages and SNMP Traps are transmitted when MIB changes occur. Default: 5 seconds. Range: 1-3600 seconds.
5. Click **OK**.

STEP 3 | Create an LLDP profile.

For descriptions of the optional TLVs, see Supported TLVs in LLDP.

1. Select **Network** > **Network Profiles** > **LLDP Profile** and **Add** a **Name** for the LLDP profile.
2. For **Mode**, select **transmit-receive** (default), **transmit-only**, or **receive-only**.
3. Select **SNMP Syslog Notification** to enable SNMP notifications and syslog messages. If enabled, the global **Notification Interval** is used. The firewall will send both an SNMP trap and a syslog event as configured in the **Device** > **Log Settings** > **System** > **SNMP Trap Profile** and **Syslog Profile**.
4. For Optional TLVs, select the TLVs you want transmitted:
   - **Port Description**
   - **System Name**
   - **System Description**
   - **System Capabilities**
5. (Optional) Select **Management Address** to add one or more management addresses and **Add** a **Name**.
6. Select the **Interface** from which to obtain the management address. At least one management address is required if **Management Address** TLV is enabled. If no management IP address is configured, the system uses the MAC address of the transmitting interface as the management address TLV.
7. Select **IPv4** or **IPv6**, and in the adjacent field, select an IP address from the drop-down (which lists the addresses configured on the selected interface), or enter an address.
8. Click **OK**.
9. Up to four management addresses are allowed. If you specify more than one **Management Address**, they will be sent in the order they are specified, starting at the top of the list. To change the order of the addresses, select an address and use the **Move Up** or **Move Down** buttons.
10. Click **OK**.

STEP 4 | Assign an LLDP profile to an interface.

1. Select **Network** > **Interfaces** and select the interface where you will assign an LLDP profile.
2. Select **Advanced** > **LLDP**.
3. Select **Enable LLDP** to assign an LLDP profile to the interface.
4. For **Profile**, select the profile you created. Selecting **None** enables LLDP with basic functionality: sends the three mandatory TLVs and enables **transmit-receive** mode.

   If you want to create a new profile, click **LLDP Profile** and follow the instructions steps above.
5. Click **OK**.

STEP 5 | Commit your changes.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

View LLDP Settings and Status

Perform the following procedure to view LLDP settings and status.

STEP 1 | View LLDP global settings.

Select **Network** > **LLDP**.

On the LLDP General screen, **Enable** indicates whether LLDP is enabled or not.

- If LLDP is enabled, the configured global settings (Transmit Interval, Transmit Delay, Hold Time Multiple, and Notification Interval) are displayed.
- If LLDP is not enabled, the default values of the global settings are displayed.

For descriptions of these values, see second step in Configure LLDP.

STEP 2 | View the LLDP status information.

1. Select the **Status** tab.
2. (Optional) Enter a filter to restrict the information that is displayed.

   Interface Information:

   - **Interface**—Name of the interfaces that have LLDP profiles assigned to them.
   - **LLDP**—LLDP status: enabled or disabled.
   - **Mode**— LLDP mode of the interface: Tx/Rx, Tx Only, or Rx Only.
   - **Profile**—Name of the profile assigned to the interface.

   Transmission Information:

   - **Total Transmitted**—Count of LLDPDUs transmitted out the interface.
   - **Dropped Transmit**—Count of LLDPDUs that were not transmitted out the interface because of an error. For example, a length error when the system is constructing an LLDPDU for transmission.

   Received Information:

   - **Total Received**—Count of LLDP frames received on the interface.
   - **Dropped TLV**—Count of LLDP frames discarded upon receipt.
   - **Errors**—Count of TLVs that were received on the interface and contained errors. Types of TLV errors include: one or more mandatory TLVs missing, out of order, containing out-of-range information, or length error.
   - **Unrecognized**—Count of TLVs received on the interface that are not recognized by the LLDP local agent. For example, the TLV type is in the reserved TLV range.
   - **Aged Out**—Count of items deleted from the Receive MIB due to proper TTL expiration.

STEP 3 | View summary LLDP information for each neighbor seen on an interface.

1. Select the **Peers** tab.
2. (Optional) Enter a filter to restrict the information being displayed.

   Local Interface—Interface on the firewall that detected the neighboring device.

   Remote Chassis ID—Chassis ID of the peer. The MAC address will be used.

   Port ID—Port ID of the peer.

   Name—Name of peer.

   More info—Provides the following remote peer details, which are based on the Mandatory and Optional TLVs:

   - Chassis Type: MAC address.
   - MAC Address: MAC address of the peer.

© 2018 Palo Alto Networks, Inc.

- System Name: Name of the peer.
- System Description: Description of the peer.
- Port Description: Port description of the peer.
- Port Type: Interface name.
- Port ID: The firewall uses the interface's ifname.
- System Capabilities: Capabilities of the system. O=Other, P=Repeater, B=Bridge, W=Wireless-LAN, R=Router, T=Telephone
- Enabled Capabilities: Capabilities enabled on the peer.
- Management Address: Management address of the peer.

Clear LLDP Statistics

You can clear LLDP statistics for specific interfaces.

Clear LLDP statistics for specific interfaces.

1. Select **Network** > **LLDP** > **Status** and in the left hand column, select one or more interfaces for which you want to clear LLDP statistics.
2. Click **Clear LLDP Statistics** at the bottom of the screen.

© 2018 Palo Alto Networks, Inc.

BFD

The firewall supports Bidirectional Forwarding Detection (BFD) (RFC 5880), a protocol that recognizes a failure in the bidirectional path between two routing peers. BFD failure detection is extremely fast, providing for a faster failover than can be achieved by link monitoring or frequent dynamic routing health checks, such as Hello packets or heartbeats. Mission-critical data centers and networks that require high availability and extremely fast failover need the extremely fast failure detection that BFD provides.

- BFD Overview
- Configure BFD
- Reference: BFD Details

BFD Overview

When you enable BFD, BFD establishes a session from one endpoint (the firewall) to its BFD peer at the endpoint of a link using a three-way handshake. Control packets perform the handshake and negotiate the parameters configured in the BFD profile, including the minimum intervals at which the peers can send and receive control packets. BFD control packets for both IPv4 and IPv6 are transmitted over UDP port 3784. BFD control packets for multihop support are transmitted over UDP port 4784. BFD control packets transmitted over either port are encapsulated in the UDP packets.

After the BFD session is established, the Palo Alto Networks implementation of BFD operates in asynchronous mode, meaning both endpoints send each other control packets (which function like Hello packets) at the negotiated interval. If a peer does not receive a control packet within the detection time (calculated as the negotiated transmit interval multiplied by a Detection Time Multiplier), the peer considers the session down. (The firewall does not support demand mode, in which control packets are sent only if necessary rather than periodically.)

When you enable BFD for a static route and a BFD session between the firewall and the BFD peer fails, the firewall removes the failed route from the RIB and FIB tables and allows an alternate path with a lower priority to take over. When you enable BFD for a routing protocol, BFD notifies the routing protocol to switch to an alternate path to the peer. Thus, the firewall and BFD peer reconverge on a new path.

A BFD profile allows you to Configure BFD settings and apply them to one or more routing protocols or static routes on the firewall. If you enable BFD without configuring a profile, the firewall uses its default BFD profile (with all of the default settings). You cannot change the default BFD profile.

When an interface is running multiple protocols that use different BFD profiles, BFD uses the profile having the lowest **Desired Minimum Tx Interval**. See BFD for Dynamic Routing Protocols.

Active/passive HA peers synchronize BFD configurations and sessions; active/active HA peers do not.

BFD is standardized in RFC 5880. PAN-OS does not support all components of RFC 5880; see Non-Supported RFC Components of BFD.

PAN-OS also supports RFC 5881, http://www.rfc-editor.org/rfc/rfc5881.txt. In this case, BFD tracks a single hop between two systems that use IPv4 or IPv6, so the two systems are directly connected to each other. BFD also tracks multiple hops from peers connected by BGP. PAN-OS follows BFD encapsulation as described in RFC 5883, https://www.rfc-editor.org/rfc/rfc5883.txt. However, PAN-OS does not support authentication.

- BFD Model, Interface, and Client Support
- Non-Supported RFC Components of BFD
- BFD for Static Routes
- BFD for Dynamic Routing Protocols

© 2018 Palo Alto Networks, Inc.

BFD Model, Interface, and Client Support

The following firewall models do not support BFD: PA-800 Series, PA-500, PA-220, PA-200, and VM-50 firewalls. The models that do support BFD support a maximum number of BFD sessions, as listed in the Product Selection tool.

BFD runs on physical Ethernet, Aggregated Ethernet (AE), VLAN, and tunnel interfaces (site-to-site VPN and LSVPN), and on Layer 3 subinterfaces.

Supported BFD clients are:

- Static routes (IPv4 and IPv6) consisting of a single hop
- OSPFv2 and OSPFv3 (interface types include broadcast, point-to point, and point-to-multipoint)
- BGP IPv4 (IBGP, EBGP) consisting of a single hop or multiple hops
- RIP (single hop)

Non-Supported RFC Components of BFD

- Demand mode
- Authentication
- Sending or receiving Echo packets; however, the firewall will pass Echo packets that arrive on a virtual wire or tap interface. (BFD Echo packets have the same IP address for the source and destination.)
- Poll sequences
- Congestion control

BFD for Static Routes

To use BFD on a static route, both the firewall and the peer at the opposite end of the static route must support BFD sessions. A static route can have a BFD profile only if the **Next Hop** type is **IP Address**.

If an interface is configured with more than one static route to a peer (the BFD session has the same source IP address and same destination IP address), a single BFD session automatically handles the multiple static routes. This behavior reduces BFD sessions. If the static routes have different BFD profiles, the profile with the smallest **Desired Minimum Tx Interval** takes effect.

In a deployment where you want to configure BFD for a static route on a DHCP or PPPoE client interface, you must perform two commits. Enabling BFD for a static route requires that the **Next Hop** type must be **IP Address**. But at the time of a DHCP or PPPoE interface commit, the interface IP address and next hop IP address (default gateway) are unknown.

You must first enable a DHCP or PPPoE client for the interface, perform a commit, and wait for the DHCP or PPPoE server to send the firewall the client IP address and default gateway IP address. Then you can configure the static route (using the default gateway address of the DHCP or PPPoE client as the next hop), enable BFD, and perform a second commit.

BFD for Dynamic Routing Protocols

In addition to BFD for static routes, the firewall supports BFD for the BGP, OSPF, and RIP routing protocols.



The Palo Alto Networks implementation of multihop BFD follows the encapsulation portion of RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths but does not support authentication. A workaround is to configure BFD in a VPN tunnel for BGP. The VPN tunnel can provide authentication without the duplication of BFD authentication.

When you enable BFD for OSPFv2 or OSPFv3 broadcast interfaces, OSPF establishes a BFD session only with its Designated Router (DR) and Backup Designated Router (BDR). On point-to-point interfaces, OSPF

© 2018 Palo Alto Networks, Inc.

establishes a BFD session with the direct neighbor. On point-to-multipoint interfaces, OSPF establishes a BFD session with each peer.

The firewall does not support BFD on an OSPF or OSPFv3 virtual link.

Each routing protocol can have independent BFD sessions on an interface. Alternatively, two or more routing protocols (BGP, OSPF, and RIP) can share a common BFD session for an interface.

When you enable BFD for multiple protocols on the same interface, and the source IP address and destination IP address for the protocols are also the same, the protocols share a single BFD session, thus reducing both dataplane overhead (CPU) and traffic load on the interface. If you configure different BFD profiles for these protocols, only one BFD profile is used: the one that has the lowest **Desired Minimum Tx Interval**. If the profiles have the same **Desired Minimum Tx Interval**, the profile used by the first created session takes effect. In the case where a static route and OSPF share the same session, because a static session is created right after a commit, while OSPF waits until an adjacency is up, the profile of the static route takes effect.

The benefit of using a single BFD session in these cases is that this behavior uses resources more efficiently. The firewall can use the saved resources to support more BFD sessions on different interfaces or support BFD for different source IP and destination IP address pairs.

IPv4 and IPv6 on the same interface always create different BFD sessions, even though they can use the same BFD profile.



*If you implement both BFD for BGP and HA path monitoring, Palo Alto Networks recommends you not implement BGP Graceful Restart. When the BFD peer's interface fails and path monitoring fails, BFD **can** remove the affected routes from the routing table and synchronize this change to the passive HA firewall before Graceful Restart can take effect. If you decide to implement BFD for BGP, Graceful Restart for BGP, and HA path monitoring, you should configure BFD with a larger Desired Minimum Tx Interval and larger Detection Time Multiplier than the default values.*

Configure BFD

After you read the BFD Overview, which includes firewall models and interfaces supported, perform the following before configuring BFD:

- Configure one or more Virtual Routers.
- Configure one or more Static Routes if you are applying BFD to static routes.
- Configure a routing protocol (BGP, OSPF, OSPFv3, or RIP) if you are applying BFD to a routing protocol.



The effectiveness of your BFD implementation depends on a variety of factors, such as traffic loads, network conditions, how aggressive your BFD settings are, and how busy the dataplane is.

STEP 1 | Create a BFD profile.



If you change a setting in a BFD profile that an existing BFD session is using and you commit the change, before the firewall deletes that BFD session and recreates it with the new setting, the firewall sends a BFD packet with the local state set to `admin down`. The peer device may or may not flap the routing protocol or static route, depending on the peer's implementation of RFC 5882, Section 3.2.

1. Select **Network** > **Network Profiles** > **BFD Profile** and **Add** a **Name** for the BFD profile. The name is case-sensitive and must be unique on the firewall. Use only letters, numbers, spaces, hyphens, and underscores.

© 2018 Palo Alto Networks, Inc.

2. Select the **Mode** in which BFD operates:
   - **Active**—BFD initiates sending control packets to peer (default). At least one of the BFD peers must be Active; both can be Active.
   - **Passive**—BFD waits for peer to send control packets and responds as required.

STEP 2 | Configure BFD intervals.

1. Enter the **Desired Minimum Tx Interval (ms)**. This is the minimum interval, in milliseconds, at which you want the BFD protocol (referred to as BFD) to send BFD control packets; you are thus negotiating the transmit interval with the peer. Minimum on PA-7000, PA-5200 Series, and PA-5000 Series firewalls is 50; minimum on PA-3000 Series firewall is 100; minimum on VM-Series firewall is 200. Maximum is 2,000; default is 1,000.

   The recommendation is to set the Desired Minimum Tx Interval on a PA-7000 or PA-5000 Series firewall to 100 or greater; a value less than 100 is at risk of causing BFD flaps.

   If you have multiple routing protocols that use different BFD profiles on the same interface, configure the BFD profiles with the same Desired Minimum Tx Interval.

2. Enter the **Required Minimum Rx Interval (ms)**. This is the minimum interval, in milliseconds, at which BFD can receive BFD control packets. Minimum on PA-7000, PA-5200 Series, and PA-5000 Series firewalls is 50; minimum on PA-3000 Series firewall is 100; minimum on VM-Series firewall is 200. Maximum is 2,000; default is 1,000.

   The recommendation is to set the Required Minimum Rx Interval on a PA-7000 or PA-5000 Series firewall to 100 or greater; a value less than 100 is at risk of causing BFD flaps.

STEP 3 | Configure the BFD Detection Time Multiplier.

Enter the **Detection Time Multiplier**. The transmit interval (negotiated from the **Desired Minimum Tx Interval**) multiplied by the **Detection Time Multiplier** equals the detection time. If BFD does not receive a BFD control packet from its peer before the detection time expires, a failure has occurred. Range is 2-50; default is 3.

For example, a transmit interval of 300 ms x 3 (Detection Time Multiplier) = 900 ms detection time.

When configuring a BFD profile, take into consideration that the firewall is a session-based device typically at the edge of a network or data center and may have slower links than a dedicated router. Therefore, the firewall likely needs a longer interval and a higher multiplier than the fastest settings allowed. A detection time that is too short can cause false failure detections when the issue is really just traffic congestion.

STEP 4 | Configure the BFD hold time.

Enter the **Hold Time (ms)**. This is the delay, in milliseconds, after a link comes up before BFD transmits BFD control packets. **Hold Time** applies to BFD Active mode only. If BFD receives BFD control packets during the **Hold Time**, it ignores them. Range is 0-120000. The default is 0, which means no transmit **Hold Time** is used; BFD sends and receives BFD control packets immediately after the link is established.

STEP 5 | (Optional—For a BGP IPv4 implementation only) Configure hop-related settings for the BFD profile.

1. Select **Multihop** to enable BFD over BGP multihop.
2. Enter the **Minimum Rx TTL**.This is the minimum Time-to-Live value (number of hops) BFD will accept (receive) in a BFD control packet when BGP supports multihop BFD. (Range is 1-254; there is no default).

© 2018 Palo Alto Networks, Inc.

The firewall drops the packet if it receives a smaller TTL than its configured **Minimum Rx TTL**. For example, if the peer is 5 hops away, and the peer transmits a BFD packet with a TTL of 100 to the firewall, and if the **Minimum Rx TTL** for the firewall is set to 96 or higher, the firewall drops the packet.

STEP 6 | Save the BFD profile.

Click **OK**.

STEP 7 | (Optional) Enable BFD for a static route.

Both the firewall and the peer at the opposite end of the static route must support BFD sessions.

1. Select **Network** > **Virtual Routers** and select the virtual router where the static route is configured.
2. Select the **Static Routes** tab.
3. Select the **IPv4** or **IPv6** tab.
4. Select the static route where you want to apply BFD.
5. Select an **Interface** (even if you are using a DHCP address). The **Interface** setting cannot be **None**.
6. For **Next Hop**, select **IP Address** and enter the IP address if not already specified.
7. For **BFD Profile**, select one of the following:
   - **default**—Uses only default settings.
   - A BFD profile you configured—See Create a BFD profile.
   - **New BFD Profile**—Allows you to Create a BFD profile.

   Selecting None (Disable BFD) disables BFD for this static route.
8. Click **OK**.

A BFD column on the **IPv4** or **IPv6** tab indicates the BFD profile configured for the static route.

STEP 8 | (Optional) Enable BFD for all BGP interfaces or for a single BGP peer.

If you enable or disable BFD globally, all interfaces running BGP will be taken down and brought back up with the BFD function. This can disrupt all BGP traffic. When you enable BFD on the interface, the firewall stops the BGP connection to the peer to program BFD on the interface. The peer device sees the BGP connection drop, which can result in a reconvergence. Enable BFD for BGP interfaces during an off-peak time when a reconvergence will not impact production traffic.

If you implement both BFD for BGP and HA path monitoring, Palo Alto Networks recommends you not implement BGP Graceful Restart. When the BFD peer's interface fails and path monitoring fails, BFD can remove the affected routes from the routing table and synchronize this change to the passive HA firewall before Graceful Restart can take effect. If you decide to implement BFD for BGP, Graceful Restart for BGP, and HA path monitoring, you should configure BFD with a larger Desired Minimum Tx Interval and larger Detection Time Multiplier than the default values.

1. Select **Network** > **Virtual Routers** and select the virtual router where BGP is configured.
2. Select the **BGP** tab.
3. (Optional) To apply BFD to all BGP interfaces on the virtual router, in the **BFD** drop-down, select one of the following and click **OK**:
   - **default**—Uses only default settings.
   - A BFD profile you configured—See Create a BFD profile.
   - **New BFD Profile**—Allows you to Create a BFD profile.

© 2018 Palo Alto Networks, Inc.

Selecting None (Disable BFD) disables BFD for all BGP interfaces on the virtual router; you cannot enable BFD for a single BGP interface.

4. (Optional) To enable BFD for a single BGP peer interface (thereby overriding the **BFD** setting for BGP as long as it is not disabled), perform the following tasks:
   1. Select the **Peer Group** tab.
   2. Select a peer group.
   3. Select a peer.
   4. In the **BFD** drop-down, select one of the following:

      default—Uses only default settings.

      Inherit-vr-global-setting (default)—The BGP peer inherits the BFD profile that you selected globally for BGP for the virtual router.

      A BFD profile you configured—See Create a BFD profile.

      Selecting Disable BFD disables BFD for the BGP peer.
5. Click **OK**.
6. Click **OK**.

A BFD column on the BGP - Peer Group/Peer list indicates the BFD profile configured for the interface.

STEP 9 | (Optional) Enable BFD for OSPF or OSPFv3 globally or for an OSPF interface.

1. Select **Network** > **Virtual Routers** and select the virtual router where OSPF or OSPFv3 is configured.
2. Select the **OSPF** or **OSPFv3** tab.
3. (Optional) In the **BFD** drop-down, select one of the following to enable BFD for all OSPF or OSPFv3 interfaces and click **OK**:
   - **default**—Uses only default settings.
   - A BFD profile you configured—See Create a BFD profile.
   - **New BFD Profile**—Allows you to Create a BFD profile.

   Selecting None (Disable BFD) disables BFD for all OSPF interfaces on the virtual router; you cannot enable BFD for a single OSPF interface.
4. (Optional) To enable BFD on a single OSPF peer interface (and thereby override the **BFD** setting for OSPF, as long as it is not disabled), perform the following tasks:
   1. Select the **Areas** tab and select an area.
   2. On the **Interface** tab, select an interface.
   3. In the **BFD** drop-down, select one of the following to configure BFD for the specified OSPF peer:

      default—Uses only default settings.

      Inherit-vr-global-setting (default)—OSPF peer inherits the **BFD** setting for OSPF or OSPFv3 for the virtual router.

      A BFD profile you configured—See Create a BFD profile.

      Selecting Disable BFD disables BFD for the OSPF or OSPFv3 interface.
   4. Click **OK**.
5. Click **OK**.

A BFD column on the OSPF **Interface** tab indicates the BFD profile configured for the interface.

© 2018 Palo Alto Networks, Inc.

STEP 10 | (Optional) Enable BFD for RIP globally or for a single RIP interface.

1. Select **Network** > **Virtual Routers** and select the virtual router where RIP is configured.
2. Select the **RIP** tab.
3. (Optional) In the **BFD** drop-down, select one of the following to enable BFD for all RIP interfaces on the virtual router and click **OK**:
   - **default**—Uses only default settings.
   - A BFD profile you configured—See Create a BFD profile.
   - **New BFD Profile**—Allows you to Create a BFD profile.

   Selecting None (Disable BFD) disables BFD for all RIP interfaces on the virtual router; you cannot enable BFD for a single RIP interface.
4. (Optional) To enable BFD for a single RIP interface (and thereby override the **BFD** setting for RIP, as long as it is not disabled), perform the following tasks:
   1. Select the **Interfaces** tab and select an interface.
   2. In the **BFD** drop-down, select one of the following:

      default—Uses only default settings).

      Inherit-vr-global-setting (default)—RIP interface inherits the BFD profile that you selected for RIP globally for the virtual router.

      A BFD profile you configured—See Create a BFD profile.

      Selecting None (Disable BFD) disables BFD for the RIP interface.
   3. Click **OK**.
5. Click **OK**.

The BFD column on the **Interface** tab indicates the BFD profile configured for the interface.

STEP 11 | Commit the configuration.

Click **Commit**.

STEP 12 | View BFD summary and details.

1. Select **Network** > **Virtual Routers**, find the virtual router you are interested in, and click **More Runtime Stats**.
2. Select the **BFD Summary Information** tab to see summary information, such as BFD state and run-time statistics.
3. (Optional) Select **details** in the row of the interface you are interested in to view Reference: BFD Details.

STEP 13 | Monitor BFD profiles referenced by a routing configuration; monitor BFD statistics, status, and state.

Use the following CLI operational commands:

- **show routing bfd active-profile [***<name>***]**
- **show routing bfd details [interface***<name>***][local-ip***<ip>***][multihop][peer-ip** *<ip>***][session-id][virtual-router***<name>***]**
- **show routing bfd drop-counters session-id** *<session-id>*
- **show counter global | match bfd**

STEP 14 | (Optional) Clear BFD transmit, receive, and drop counters.

© 2018 Palo Alto Networks, Inc.

```
clear routing bfd counters session-id all | <1-1024>
```

STEP 15 | (Optional) Clear BFD sessions for debugging.

```
clear routing bfd session-state session-id all | <1-1024>
```

Reference: BFD Details

To see the following BFD information for a virtual router, you can refer Step View BFD summary and details.

Name	Value (Example)	Description
Session ID	1	ID number of the BFD session.
Interface	ethernet1/12	Interface you selected where BFD is running.
Protocol	STATIC(IPV4) OSPF	Static route (IP address family of static route) and/or dynamic routing protocol that is running BFD on the interface.
Local IP Address	10.55.55.2	IP address of interface.
Neighbor IP Address	10.55.55.1	IP address of BFD neighbor.
BFD Profile	default *(This BFD session has multiple BFD profiles. Lowest 'Desired Minimum Tx Interval (ms)' is used to select the effective profile.)	Name of BFD profile applied to the interface. Because the sample interface has both a static route and OSPF running BFD with different profiles, the firewall uses the profile with the lowest **Desired Minimum Tx Interval**. In this example, the profile used is the default profile.
State (local/remote)	up/up	BFD states of the local and remote BFD peers. Possible states are admin down, down, init, and up.
Up Time	2h 36m 21s 419ms	Length of time BFD has been up (hours, minutes, seconds, and milliseconds).
Discriminator (local/ remote)	1391591427/1	Discriminators for local and remote BFD peers.
Mode	Active	Mode in which BFD is configured on the interface: Active or Passive.
Demand Mode	Disabled	PAN-OS does not support BFD Demand Mode, so it is always in Disabled state.
Multihop	Disabled	BFD multihop: Enabled or Disabled.

© 2018 Palo Alto Networks, Inc.

Name	Value (Example)	Description
Multihop TTL		TTL of multihop; range is 1-254. Field is empty if Multihop is disabled.
Local Diag Code	0 (No Diagnostic)	Diagnostic codes indicating the reason for the local system's last change in state: 0—No Diagnostic 1—Control Detection Time Expired 2—Echo Function Failed 3—Neighbor Signaled Session Down 4—Forwarding Plane Reset 5—Path Down 6—Concatenated Path Down 7—Administratively Down 8—Reverse Concatenated Path Down
Last Received Remote Diag Code	0 (No Diagnostic)	Diagnostic code last received from BFD peer.
Transmit Hold Time	0ms	Hold time (in milliseconds) after a link comes up before BFD transmits BFD control packets. A hold time of 0ms means to transmit immediately. Range is 0-120000ms.
Received Min Rx Interval	1000ms	Minimum Rx interval received from the peer; the interval at which the BFD peer can receive control packets. Maximum is 2000ms.
Negotiated Transmit Interval	1000ms	Transmit interval (in milliseconds) that the BFD peers have agreed to send BFD control packets to each other. Maximum is 2000ms.
Received Multiplier	3	Detection time multiplier value received from the BFD peer. The Transmit Time multiplied by the Multiplier equals the detection time. If BFD does not receive a BFD control packet from its peer before the detection time expires, a failure has occurred. Range is 2-50.
Detect Time (exceeded)	3000ms (0)	Calculated detection time (Negotiated Transmit Interval multiplied by Multiplier) and the number of milliseconds the detection time is exceeded.
Tx Control Packets (last)	9383 (420ms ago)	Number of BFD control packets transmitted (and length of time since BFD transmitted the most recent control packet).
Rx Control Packets (last)	9384 (407ms ago)	Number of BFD control packets received (and length of time since BFD received the most recent control packet).

© 2018 Palo Alto Networks, Inc.

Name	Value (Example)	Description
Agent Data Plane	Slot 1 - DP 0	On PA-7000 Series firewalls, the dataplane CPU that is assigned to handle packets for this BFD session.
Errors	0	Number of BFD errors.
Last Packet Causing State Change		
Version	1	BFD version.
Poll Bit	0	BFD poll bit; 0 indicates not set.
Desired Min Tx Interval	1000ms	Desired minimum transmit interval of last packet causing state change.
Required Min Rx Interval	1000ms	Required minimum receive interval of last packet causing state change.
Detect Multiplier	3	Detect Multiplier of last packet causing state change.
My Discriminator	1	Remote discriminator. A discriminator is a unique, nonzero value the peers use to distinguish multiple BFD sessions between them.
Your Discriminator	1391591427	Local discriminator. A discriminator is a unique, nonzero value the peers use to distinguish multiple BFD sessions between them.
Diagnostic Code	0 (No Diagnostic)	Diagnostic code of last packet causing state change.
Length	24	Length of BFD control packet in bytes.
Demand Bit	0	PAN-OS does not support BFD Demand mode, so Demand Bit is always set to 0 (disabled).
Final Bit	0	PAN-OS does not support the Poll Sequence, so Final Bit is always set to 0 (disabled).
Multipoint Bit	0	This bit is reserved for future point-to-multipoint extensions to BFD. It must be zero on both transmit and receipt.
Control Plane Independent Bit	1	• If set to 1, the transmitting system's BFD implementation does not share fate with its control plane (i.e., BFD is implemented in the forwarding plane and can continue to function through disruptions in the control plane). In PAN-OS, this bit is always set to 1. • If set to 0, the transmitting system's BFD implementation shares fate with its control plane.

© 2018 Palo Alto Networks, Inc.

Name	Value (Example)	Description
Authentication Present Bit	0	PAN-OS does not support BFD Authentication, so the Authentication Present Bit is always set to 0.
Required Min Echo Rx Interval	0ms	PAN-OS does not support the BFD Echo function, so this will always be 0ms.

© 2018 Palo Alto Networks, Inc.

Session Settings and Timeouts

This section describes the global settings that affect TCP, UDP, and ICMPv6 sessions, in addition to IPv6, NAT64, NAT oversubscription, jumbo frame size, MTU, accelerated aging, and Captive Portal authentication. There is also a setting (Rematch Sessions) that allows you to apply newly configured security policies to sessions that are already in progress.

The first few topics below provide brief summaries of the Transport Layer of the OSI model, TCP, UDP, and ICMP. For more information about the protocols, refer to their respective RFCs. The remaining topics describe the session timeouts and settings.

- Transport Layer Sessions
- TCP
- UDP
- ICMP
- Control Specific ICMP or ICMPv6 Types and Codes
- Configure Session Timeouts
- Session Distribution Policies
- Configure Session Settings
- Prevent TCP Split Handshake Session Establishment

Transport Layer Sessions

A network session is an exchange of messages that occurs between two or more communication devices, lasting for some period of time. A session is established and is torn down when the session ends. Different types of sessions occur at three layers of the OSI model: the Transport layer, the Session layer, and the Application layer.

The Transport Layer operates at Layer 4 of the OSI model, providing reliable or unreliable, end-to-end delivery and flow control of data. Internet protocols that implement sessions at the Transport layer include Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

TCP

Transmission Control Protocol (TCP) (RFC 793) is one of the main protocols in the Internet Protocol (IP) suite, and is so prevalent that it is frequently referenced together with IP as *TCP/IP*. TCP is considered a reliable transport protocol because it provides error-checking while transmitting and receiving segments, acknowledges segments received, and reorders segments that arrive in the wrong order. TCP also requests and provides retransmission of segments that were dropped. TCP is stateful and connection-oriented, meaning a connection between the sender and receiver is established for the duration of the session. TCP provides flow control of packets, so it can handle congestion over networks.

TCP performs a handshake during session setup to initiate and acknowledge a session. After the data is transferred, the session is closed in an orderly manner, where each side transmits a FIN packet and acknowledges it with an ACK packet. The handshake that initiates the TCP session is often a three-way handshake (an exchange of three messages) between the initiator and the listener, or it could be a variation, such as a four-way or five-way split handshake or a simultaneous open. The TCP Split Handshake Drop on page 961 explains how to Prevent TCP Split Handshake Session Establishment on page 973.

Applications that use TCP as their transport protocol include Hypertext Transfer Protocol (HTTP), HTTP Secure (HTTPS), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), Telnet, Post Office Protocol version 3 (POP3), Internet Message Access Protocol (IMAP), and Secure Shell (SSH).

The following topics describe details of the PAN-OS implementation of TCP.

© 2018 Palo Alto Networks, Inc.

- TCP Half Closed and TCP Time Wait Timers on page 960
- Unverified RST Timer on page 961
- TCP Split Handshake Drop on page 961
- Maximum Segment Size (MSS) on page 962

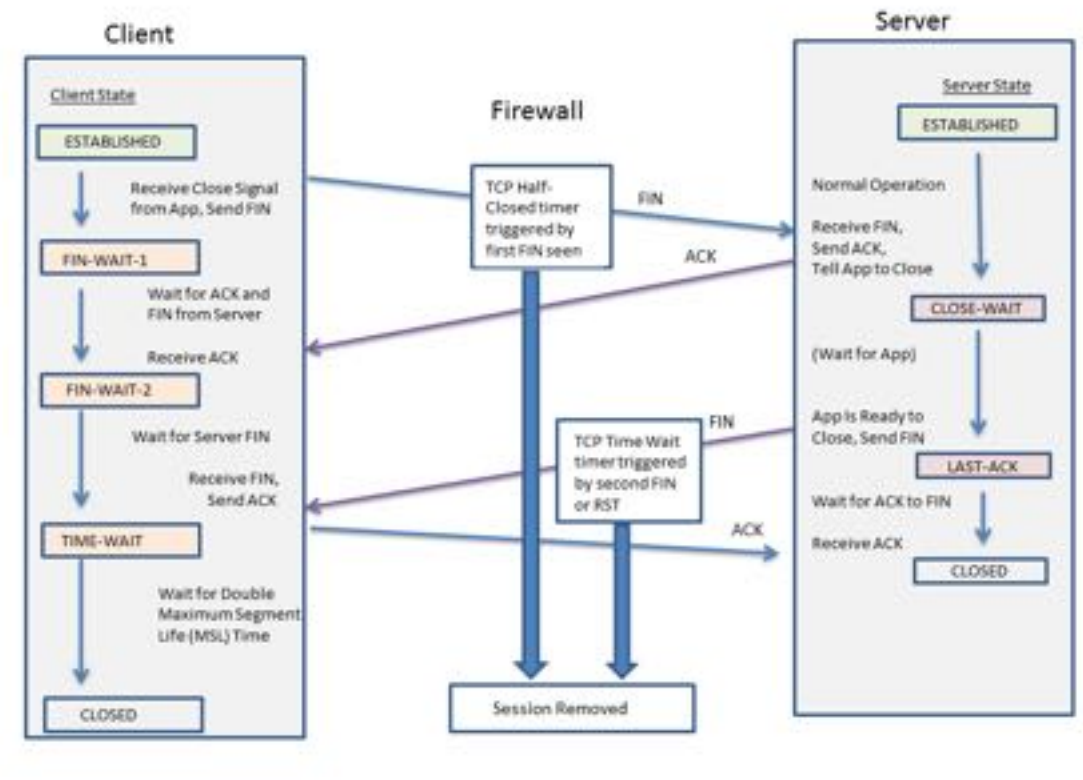
You can use Zone Protection Profiles on page 1069 on the firewall to configure packet-based attack protection and thereby drop IP, TCP, and IPv6 packets with undesirable characteristics or strip undesirable options from packets before allowing them into the zone. You can also configure flood protection, specifying the rate of SYN connections per second (not matching an existing session) that trigger an alarm, cause the firewall to randomly drop SYN packets or use SYN cookies, and cause the firewall to drop SYN packets that exceed the maximum rate.

TCP Half Closed and TCP Time Wait Timers

The TCP connection termination procedure uses a TCP Half Closed timer, which is triggered by the first FIN the firewall sees for a session. The timer is named TCP Half Closed because only one side of the connection has sent a FIN. A second timer, TCP Time Wait, is triggered by the second FIN or a RST.

If the firewall were to have only one timer triggered by the first FIN, a setting that was too short could prematurely close the half-closed sessions. Conversely, a setting that was too long would make the session table grow too much and possibly use up all of the sessions. Two timers allow you to have a relatively long TCP Half Closed timer and a short TCP Time Wait timer, thereby quickly aging fully closed sessions and controlling the size of the session table.

The following figure illustrates when the firewall's two timers are triggered during the TCP connection termination procedure.



The TCP Time Wait timer should be set to a value less than the TCP Half Closed timer for the following reasons:

© 2018 Palo Alto Networks, Inc.

- The longer time allowed after the first FIN is seen gives the opposite side of the connection time to fully close the session.
- The shorter Time Wait time is because there is no need for the session to remain open for a long time after the second FIN or a RST is seen. A shorter Time Wait time frees up resources sooner, yet still allows time for the firewall to see the final ACK and possible retransmission of other datagrams.

If you configure a TCP Time Wait timer to a value greater than the TCP Half Closed timer, the commit will be accepted, but in practice the TCP Time Wait timer will not exceed the TCP Half Closed value.

The timers can be set globally or per application. The global settings are used for all applications by default. If you configure TCP wait timers at the application level, they override the global settings.

Unverified RST Timer

If the firewall receives a Reset (RST) packet that cannot be verified (because it has an unexpected sequence number within the TCP window or it is from an asymmetric path), the Unverified RST timer controls the aging out of the session. It defaults to 30 seconds; the range is 1-600 seconds. The Unverified RST timer provides an additional security measure, explained in the second bullet below.

A RST packet will have one of three possible outcomes:

- A RST packet that falls outside the TCP window is dropped.
- A RST packet that falls inside the TCP window but does not have the exact expected sequence number is unverified and subject to the Unverified RST timer setting. This behavior helps prevent denial of service (DoS) attacks where the attack tries to disrupt existing sessions by sending random RST packets to the firewall.
- A RST packet that falls within the TCP window and has the exact expected sequence number is subject to the TCP Time Wait timer setting.

TCP Split Handshake Drop

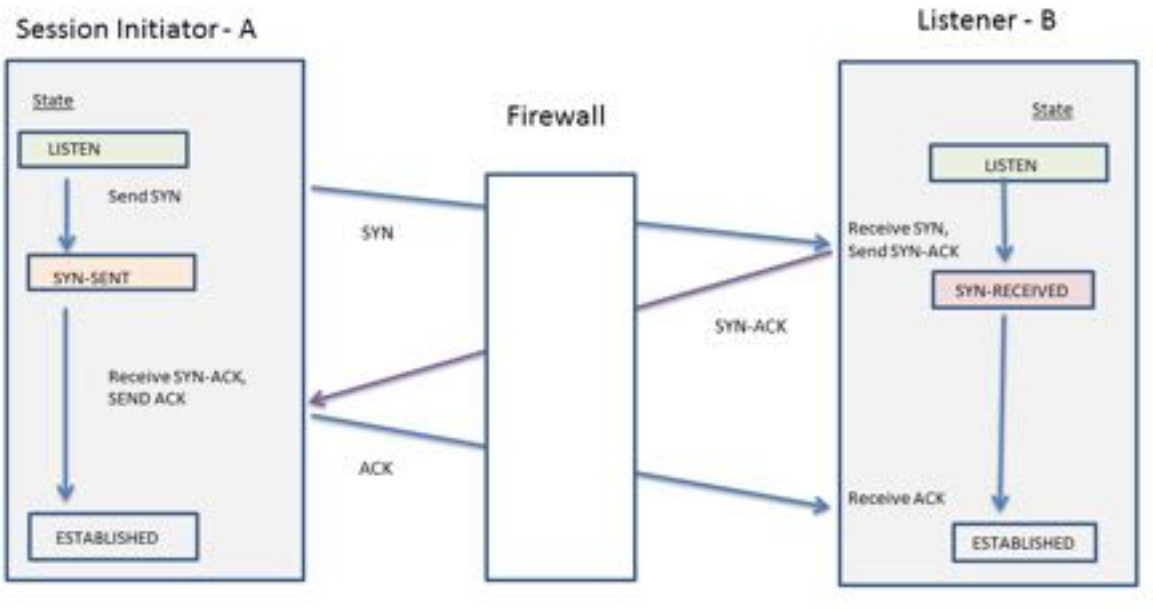
The **Split Handshake** option in a Zone Protection profile will prevent a TCP session from being established if the session establishment procedure does not use the well-known three-way handshake, but instead uses a variation, such as a four-way or five-way split handshake or a simultaneous open.

The Palo Alto Networks next-generation firewall correctly handles sessions and all Layer 7 processes for split handshake and simultaneous open session establishment without enabling the **Split Handshake** option. Nevertheless, the **Split Handshake** option (which causes a TCP split handshake drop) is made available. When the **Split Handshake** option is configured for a Zone Protection profile and that profile is applied to a zone, TCP sessions for interfaces in that zone must be established using the standard three-way handshake; variations are not allowed.

The **Split Handshake** option is disabled by default.

The following illustrates the standard three-way handshake used to establish a TCP session with a PAN-OS firewall between the initiator (typically a client) and the listener (typically a server).

© 2018 Palo Alto Networks, Inc.



The **Split Handshake** option is configured for a Zone Protection profile that is assigned to a zone. An interface that is a member of the zone drops any synchronization (SYN) packets sent from the server, preventing the following variations of handshakes. The letter A in the figure indicates the session initiator and B indicates the listener. Each numbered segment of the handshake has an arrow indicating the direction of the segment from the sender to the receiver, and each segment indicates the control bit(s) setting.

4-Way Split Handshake (Version 1)	4-Way Split Handshake (Version 2)	Simultaneous Open	5-Way Split Handshake
1. A → B SYN	1. A → B SYN	1. A → B SYN	1. A → B SYN
2. A ← B ACK	2. A ← B SYN	2. A ← B SYN	2. A ← B ACK
3. A ← B SYN	3. A → B SYN-ACK	3. A → B SYN-ACK	3. A ← B SYN
4. A → B ACK	4. A ← B ACK	4. A ← B SYN-ACK	4. A → B SYN-ACK
			5. A ← B ACK

You can Prevent TCP Split Handshake Session Establishment.

Maximum Segment Size (MSS)

The maximum transmission unit (MTU) is a value indicating the largest number of bytes that can be transmitted in a single TCP packet. The MTU includes the length of headers, so the MTU minus the number of bytes in the headers equals the maximum segment size (MSS), which is the maximum number of data bytes that can be transmitted in a single packet.

A configurable MSS adjustment size (shown below) allows your firewall to pass traffic that has longer headers than the default setting allows. Encapsulation adds length to headers, so you would increase the MSS adjustment size to allow bytes, for example, to accommodate an MPLS header or tunneled traffic that has a VLAN tag.

© 2018 Palo Alto Networks, Inc.



If the DF (don't fragment) bit is set for a packet, it is especially helpful to have a larger MSS adjustment size and smaller MSS so that longer headers do not result in a packet length that exceeds the allowed MTU. If the DF bit were set and the MTU were exceeded, the larger packets would be dropped.

The firewall supports a configurable MSS adjustment size for IPv4 and IPv6 addresses on the following Layer 3 interface types: Ethernet, subinterfaces, Aggregated Ethernet (AE), VLAN, and loopback. The IPv6 MSS adjustment size applies only if IPv6 is enabled on the interface.

If IPv4 and IPv6 are enabled on an interface and the MSS Adjustment Size differs between the two IP address formats, the proper MSS value corresponding to the IP type is used for TCP traffic.

For IPv4 and IPv6 addresses, the firewall accommodates larger-than-expected TCP header lengths. In the case where a TCP packet has a larger header length than you planned for, the firewall chooses as the MSS adjustment size the larger of the following two values:

- The configured MSS adjustment size
- The sum of the length of the TCP header (20) + the length of IP headers in the TCP SYN

This behavior means that the firewall overrides the configured MSS adjustment size if necessary. For example, if you configure an MSS adjustment size of 42, you expect the MSS to equal 1458 (the default MTU size minus the adjustment size [1500 - 42]). However, the TCP packet has 4 extra bytes of IP options in the header, so the MSS adjustment size (20+20+4) equals 44, which is larger than the configured MSS adjustment size of 42. The resulting MSS is 1500-44=1456 bytes, smaller than you expected.

To configure the MSS adjustment size, see Tune the Maximum Segment Size (MSS) adjustment size settings for a Layer 3 interface. in Configure Session Settings.

UDP

User Datagram Protocol (UDP) (RFC 768) is another main protocol of the IP suite, and is an alternative to TCP. UDP is stateless and connectionless in that there is no handshake to set up a session, and no connection between the sender and receiver; the packets may take different routes to get to a single destination. UDP is considered an unreliable protocol because it does not provide acknowledgments, error-checking, retransmission, or reordering of datagrams. Without the overhead required to provide those features, UDP has reduced latency and is faster than TCP. UDP is referred to as a best-effort protocol because there is no mechanism or guarantee to ensure that the data will arrive at its destination.

A UDP datagram is encapsulated in an IP packet. Although UDP uses a checksum for data integrity, it performs no error checking at the network interface level. Error checking is assumed to be unnecessary or is performed by the application rather than UDP itself. UDP has no mechanism to handle flow control of packets.

UDP is often used for applications that require faster speeds and time-sensitive, real-time delivery, such as Voice over IP (VoIP), streaming audio and video, and online games. UDP is transaction-oriented, so it is also used for applications that respond to small queries from many clients, such as Domain Name System (DNS) and Trivial File Transfer Protocol (TFTP).

You can use Zone Protection Profiles on page 1069 on the firewall to configure flood protection and thereby specify the rate of UDP connections per second (not matching an existing session) that trigger an alarm, trigger the firewall to randomly drop UDP packets, and cause the firewall to drop UDP packets that exceed the maximum rate. (Although UDP is connectionless, the firewall tracks UDP datagrams in IP

© 2018 Palo Alto Networks, Inc.

packets on a session basis; therefore if the UDP packet doesn't match an existing session, it is considered a new session and it counts as a connection toward the thresholds.)

ICMP

Internet Control Message Protocol (ICMP) (RFC 792) is another one of the main protocols of the Internet Protocol suite; it operates at the Network layer of the OSI model. ICMP is used for diagnostic and control purposes, to send error messages about IP operations, or messages about requested services or the reachability of a host or router. Network utilities such as traceroute and ping are implemented by using various ICMP messages.

ICMP is a connectionless protocol that does not open or maintain actual sessions. However, the ICMP messages between two devices can be considered a session.

Palo Alto Networks firewalls support ICMPv4 and ICMPv6. You can control ICMPv4 and ICMPv6 packets in several ways:

- Create Security Policy Rules Based on ICMP and ICMPv6 Packets on page 964 and select the **icmp** or **ipv6-icmp** application in the rule.
- Control ICMPv6 Rate Limiting on page 965 when you Configure Session Settings on page 967.
- Use Zone Protection Profiles on page 1069 to configure flood protection, specifying the rate of ICMP or ICMPv6 connections per second (not matching an existing session) that trigger an alarm, trigger the firewall to randomly drop ICMP or ICMPv6 packets, and cause the firewall to drop ICMP or ICMPv6 packets that exceed the maximum rate.
- Use Zone Protection Profiles on page 1069 to configure packet based attack protection:
  - For ICMP, you can drop certain types of packets or suppress the sending of certain packets.
  - For ICMPv6 packets (Types 1, 2, 3, 4, and 137), you can specify that the firewall use the ICMP session key to match a security policy rule, which determines whether the ICMPv6 packet is allowed or not. (The firewall uses the security policy rule, overriding the default behavior of using the embedded packet to determine a session match.) When the firewall drops ICMPv6 packets that match a security policy rule, the firewall logs the details in Traffic logs.

Security Policy Rules Based on ICMP and ICMPv6 Packets

The firewall forwards ICMP or ICMPv6 packets only if a security policy rule allows the session (as the firewall does for other packet types). The firewall determines a session match in one of two ways, depending on whether the packet is an ICMP or ICMPv6 error packet or redirect packet as opposed to an ICMP or ICMPv6 informational packet:

- **ICMP Types 3, 5, 11, and 12 and ICMPv6 Types 1, 2, 3, 4, and 137**—The firewall by default looks up the embedded IP packet bytes of information from the original datagram that caused the error (the invoking packet). If the embedded packet matches an existing session, the firewall forwards or drops the ICMP or ICMPv6 packet according to the action specified in the security policy rule that matches that same session. (You can use Zone Protection Profiles on page 1069 with packet based attack protection to override this default behavior for the ICMPv6 types.)
- **Remaining ICMP or ICMPv6 Packet Types**—The firewall treats the ICMP or ICMPv6 packet as if it belongs to a new session. If a security policy rule matches the packet (which the firewall recognizes as an **icmp** or **ipv6-icmp** session), the firewall forwards or drops the packet based on the security policy rule action. Security policy counters and traffic logs reflect the actions.

  If no security policy rule matches the packet, the firewall applies its default security policy rules, which allow intrazone traffic and block interzone traffic (logging is disabled by default for these rules).



Although you can override the default rules to enable logging or change the default action, we don't recommend you change the default behavior for a specific case because

© 2018 Palo Alto Networks, Inc.

it will impact all traffic that those default rules affect. Instead, create security policy rules to control and log ICMP or ICMPv6 packets explicitly.

There are two ways to create explicit security policy rules to handle ICMP or ICMPv6 packets that are not error or redirect packets:

- **Create a security policy rule to allow (or deny) all ICMP or ICMPv6 packets**—In the security policy rule, specify the application **icmp** or **ipv6-icmp**; the firewall allows (or denies) all IP packets matching the ICMP protocol number (1) or ICMPv6 protocol number (58), respectively, through the firewall.
- **Create a custom application and a security policy rule to allow (or deny) packets from or to that application**—This more granular approach allows you to Control Specific ICMP or ICMPv6 Types and Codes on page 965.

ICMPv6 Rate Limiting

ICMPv6 rate limiting is a throttling mechanism to prevent flooding and DDoS attempts. The implementation employs an error packet rate and a token bucket, which work together to enable throttling and ensure that ICMP packets don't flood the network segments protected by the firewall.

First the global **ICMPv6 Error Packet Rate (per sec)** controls the rate at which ICMPv6 error packets are allowed through the firewall; the default is 100 packets per second; the range is 10 to 65535 packets per second. If the firewall reaches the ICMPv6 error packet rate, then the token bucket comes into play and throttling occurs, as follows.

The concept of a logical token bucket controls the rate at which ICMP messages can be transmitted. The number of tokens in the bucket is configurable, and each token represents an ICMPv6 message that can be sent. The token count is decremented each time an ICMPv6 message is sent; when the bucket reaches zero tokens, no more ICMPv6 messages can be sent until another token is added to the bucket. The default size of the token bucket is 100 tokens (packets); the range is 10 to 65535 tokens.

To change the default token bucket size or error packet rate, see the section Configure Session Settings on page 967.

Control Specific ICMP or ICMPv6 Types and Codes

Use this task to create a custom ICMP or ICMPv6 application and then create a security policy rule to allow or deny that application.

STEP 1 | Create a custom application for ICMP or ICMPv6 message types and codes.

1. Select **Object** > **Applications** and **Add** a custom application.
2. On the **Configuration** tab, enter a **Name** for the custom application and a **Description**. For example, enter the name ping6.
3. For **Category**, select **networking**.
4. For **Subcategory**, select **ip-protocol**.
5. For **Technology**, select **network-protocol**.
6. Click **OK**.
7. On the **Advanced** tab, select **ICMP Type** or **ICMPv6 Type**.
8. For **Type**, enter the number (range is 0-255) that designates the ICMP or ICMPv6 message type you want to allow or deny. For example, Echo Request message (ping) is 128.
9. If the Type includes codes, enter the **Code** number (range is 0-255) that applies to the **Type** value you want to allow or deny. Some **Type** values have Code 0 only.
10. Click **OK**.

STEP 2 | Create a Security policy rule that allows or denies the custom application you created.

© 2018 Palo Alto Networks, Inc.

Create a Security Policy Rule. On the **Application** tab, specify the name of the custom application you just created.

STEP 3 | Commit your changes.

Click **Commit**.

Configure Session Timeouts

A session timeout defines the duration of time for which PAN-OS maintains a session on the firewall after inactivity in the session. By default, when the session timeout for the protocol expires, PAN-OS closes the session. You can define a number of timeouts for TCP, UDP, and ICMP sessions in particular. The Default timeout applies to any other type of session. The timeouts are global, meaning they apply to all of the sessions of that type on the firewall.

In addition to the global settings, you can define timeouts for an individual application in the **Objects** > **Applications** tab. The firewall applies application timeouts to an application that is in established state. When configured, timeouts for an application override the global TCP or UDP session timeouts.



If you change the timers at the application level, predefined applications and shared applications will all be the same, regardless of whether they're for shared virtual systems or per virtual system. If you need an application's timers to be different for a virtual system, you must create a custom application, assign it unique timers, and then assign the custom application to a unique virtual system.

Perform the task below if you need to change default values of the global session timeout settings for TCP, UDP, ICMP, Captive Portal authentication, or other types of sessions. All values are in seconds.



The defaults are optimal values. However, you can modify these according to your network needs. Setting a value too low could cause sensitivity to minor network delays and could result in a failure to establish connections with the firewall. Setting a value too high could delay failure detection.

STEP 1 | Access the Session Settings.

Select **Device** > **Setup** > **Session** and edit the Session Timeouts.

STEP 2 | (Optional) Change miscellaneous timeouts.

- **Default**—Maximum length of time that a non-TCP/UDP or non-ICMP session can be open without a response (range is 1-15,999,999; default is 30).
- **Discard Default**—Maximum length of time that a non-TCP/UDP session remains open after PAN-OS denies a session based on security policies configured on the firewall (range is 1-15,999,999; default is 60).
- **Scan**—Maximum length of time that any session remains open after it is considered inactive; an application is regarded as inactive when it exceeds the application trickling threshold defined for the application (range is 5-30; default is 10).
- **Captive Portal**—Authentication session timeout for the Captive Portal web form. To access the requested content, the user must enter the authentication credentials in this form and be successfully authenticated (range is 1-15,999,999; default is 30).

To define other Captive Portal timeouts, such as the idle timer and the expiration time before the user must be re-authenticated, select **Device** > **User Identification** > **Captive Portal Settings**. See Configure Captive Portal

© 2018 Palo Alto Networks, Inc.

STEP 3 | (Optional) Change TCP timeouts.

- **Discard TCP**—Maximum length of time that a TCP session remains open after it is denied based on a security policy configured on the firewall. Default: 90. Range: 1-15,999,999.
- **TCP**—Maximum length of time that a TCP session remains open without a response, after a TCP session is in the Established state (after the handshake is complete and/or data is being transmitted). Default: 3,600. Range: 1-15,999,999.
- **TCP Handshake**—Maximum length of time permitted between receiving the SYN-ACK and the subsequent ACK to fully establish the session. Default: 10. Range: 1-60.
- **TCP init**—Maximum length of time permitted between receiving the SYN and SYN-ACK prior to starting the TCP handshake timer. Default: 5. Range: 1-60.
- **TCP Half Closed**—Maximum length of time between receiving the first FIN and receiving the second FIN or a RST. Default: 120. Range: 1-604,800.
- **TCP Time Wait**—Maximum length of time after receiving the second FIN or a RST. Default: 15. Range: 1-600.
- **Unverified RST**—Maximum length of time after receiving a RST that cannot be verified (the RST is within the TCP window but has an unexpected sequence number, or the RST is from an asymmetric path). Default: 30. Range: 1-600.
- See also the **Scan** timeout in the section (Optional) Change miscellaneous timeouts.

STEP 4 | (Optional) Change UDP timeouts.

- **Discard UDP**—Maximum length of time that a UDP session remains open after it is denied based on a security policy configured on the firewall. Default: 60. Range: 1-15,999,999.
- **UDP**—Maximum length of time that a UDP session remains open without a UDP response. Default: 30. Range: 1-15,999,999.
- See also the **Scan** timeout in the section (Optional) Change miscellaneous timeouts.

STEP 5 | (Optional) Change ICMP timeouts.

- **ICMP**—Maximum length of time that an ICMP session can be open without an ICMP response. Default: 6. Range: 1-15,999,999.
- See also the **Discard Default** and **Scan** timeout in the section (Optional) Change miscellaneous timeouts.

STEP 6 | Commit your changes.

Click **OK** and **Commit**.

Configure Session Settings

This topic describes various settings for sessions other than timeouts values. Perform these tasks if you need to change the default settings.

STEP 1 | Change the session settings.

Select **Device** > **Setup** > **Session** and edit the Session Settings.

STEP 2 | Specify whether to apply newly configured Security policy rules to sessions that are in progress.

Select **Rematch all sessions on config policy change** to apply newly configured Security policy rules to sessions that are already in progress. This capability is enabled by default. If you clear this check box, any policy rule changes you make apply only to sessions initiated after you commit the policy change.

© 2018 Palo Alto Networks, Inc.

For example, if a Telnet session started while an associated policy rule was configured that allowed Telnet, and you subsequently committed a policy change to deny Telnet, the firewall applies the revised policy to the current session and blocks it.

STEP 3 | Configure IPv6 settings.

- **ICMPv6 Token Bucket Size**—Default: 100 tokens. See the section ICMPv6 Rate Limiting.
- **ICMPv6 Error Packet Rate (per sec)**—Default: 100. See the section ICMPv6 Rate Limiting.
- **Enable IPv6 Firewalling**—Enables firewall capabilities for IPv6. All IPv6-based configurations are ignored if IPv6 is not enabled. Even if IPv6 is enabled for an interface, the **IPv6 Firewalling** setting must also be enabled for IPv6 to function.

STEP 4 | Enable jumbo frames and set the MTU.

1. Select **Enable Jumbo Frame** to enable jumbo frame support on Ethernet interfaces. Jumbo frames have a maximum transmission unit (MTU) of 9,216 bytes and are available on certain models.
2. Set the **Global MTU**, depending on whether or not you enabled jumbo frames:
   - If you did not enable jumbo frames, the **Global MTU** defaults to 1,500 bytes; the range is 576 to 1,500 bytes.
   - If you enabled jumbo frames, the **Global MTU** defaults to 9,192 bytes; the range is 9,192 to 9,216 bytes.

If you enable jumbo frames and you have interfaces where the MTU is not specifically configured, those interfaces will automatically inherit the jumbo frame size. Therefore, before you enable jumbo frames, if you have any interface that you do not want to have jumbo frames, you must set the MTU for that interface to 1500 bytes or another value.

STEP 5 | Tune NAT session settings.

- **NAT64 IPv6 Minimum Network MTU**—Sets the global MTU for IPv6 translated traffic. The default of 1,280 bytes is based on the standard minimum MTU for IPv6 traffic.
- **NAT Oversubscription Rate**—If NAT is configured to be Dynamic IP and Port (DIPP) translation, an oversubscription rate can be configured to multiply the number of times that the same translated IP address and port pair can be used concurrently. The rate is 1, 2, 4, or 8. The default setting is based on the firewall model.
- A rate of 1 means no oversubscription; each translated IP address and port pair can be used only once at a time.
- If the setting is **Platform Default**, user configuration of the rate is disabled and the default oversubscription rate for the model applies.

Reducing the oversubscription rate decreases the number of source device translations, but provides higher NAT rule capacities.

STEP 6 | Tune accelerated aging settings.

Select **Accelerated Aging** to enable faster aging-out of idle sessions. You can also change the threshold (%) and scaling factor:

- **Accelerated Aging Threshold**—Percentage of the session table that is full when accelerated aging begins. The default is 80%. When the session table reaches this threshold (% full), PAN-OS applies the Accelerated Aging Scaling Factor to the aging calculations for all sessions.
- **Accelerated Aging Scaling Factor**—Scaling factor used in the accelerated aging calculations. The default scaling factor is 2, meaning that the accelerated aging occurs at a rate twice as fast as the configured idle time. The configured idle time divided by 2 results in a faster timeout of one-half the

© 2018 Palo Alto Networks, Inc.

time. To calculate the session's accelerated aging, PAN-OS divides the configured idle time (for that type of session) by the scaling factor to determine a shorter timeout.

For example, if the scaling factor is 10, a session that would normally time out after 3600 seconds would time out 10 times faster (in 1/10 of the time), which is 360 seconds.

STEP 7 | Enable packet buffer protection.

1. Select **Packet Buffer Protection** to enable the firewall to take action against sessions that can overwhelm the its packet buffer and causes legitimate traffic to be dropped.
2. If you enable packet buffer protection, you can tune the thresholds and timers that dictate how the firewall responds to packet buffer abuse.
   - **Alert (%)**: When packet buffer utilization exceeds this threshold, the firewall creates a log event. The threshold is set to 50% by default and the range is 0% to 99%. If the value is set to 0%, the firewall does not create a log event.
   - **Activate (%)**: When a packet buffer utilization exceeds this threshold, the firewall applies random early drop (RED) to abusive sessions. The threshold is set to 50% by default and the range is 0% to 99%. If the value is set to 0%, the firewall does not apply RED.

   Alert events are recorded in the system log. Events for dropped traffic, discarded sessions, and blocked IP address are recorded in the threat log.

   - **Block Hold Time (sec)**: The amount of time a RED-mitigated session is allowed to continue before it is discarded. By default, the block hold time is 60 seconds. The range is 0 to 65,535 seconds. If the value is set to 0, the firewall does not discard sessions based on packet buffer protection.
   - **Block Duration (sec)**: This setting defines how long a session is discarded or an IP address is blocked. The default is 3,600 second with a range of 0 seconds to 15,999,999 seconds. If this value is set to 0, the firewall does not discard sessions or block IP addresses based on packet buffer protection.

STEP 8 | Enable buffering of multicast route setup packets.

1. Select **Multicast Route Setup Buffering** to enable the firewall to preserve the first packet in a multicast session when the multicast route or forwarding information base (FIB) entry does not yet exist for the corresponding multicast group. By default, the firewall does not buffer the first multicast packet in a new session; instead, it uses the first packet to set up the multicast route. This is expected behavior for multicast traffic. You only need to enable multicast route setup buffering if your content servers are directly connected to the firewall and your custom application cannot withstand the first packet in the session being dropped. This option is disabled by default.
2. If you enable buffering, you can also tune the **Buffer Size**, which specifies the buffer size per flow. The firewall can buffer a maximum of 5,000 packets.

   You can also tune the duration, in seconds, for which a multicast route remains in the routing table on the firewall after the session ends by configuring the multicast settings on the virtual router that handles your virtual router (set the Multicast Route Age Out Time (sec) on the Multicast > Advanced tab in the virtual router configuration.

STEP 9 | Save the session settings.

Click **OK**.

STEP 10 | Tune the Maximum Segment Size (MSS) adjustment size settings for a Layer 3 interface.

1. Select **Network** > **Interfaces**, select **Ethernet**, **VLAN**, or **Loopback**, and select a Layer 3 interface.
2. Select **Advanced** > **Other Info**.
3. Select **Adjust TCP MSS** and enter a value for one or both of the following:

© 2018 Palo Alto Networks, Inc.

- **IPv4 MSS Adjustment Size** (range is 40-300 bytes; default is 40 bytes).
- **IPv6 MSS Adjustment Size** (range is 60-300 bytes; default is 60 bytes).

4. Click **OK**.

STEP 11 | Commit your changes.

Click **Commit**.

STEP 12 | Reboot the firewall after changing the jumbo frame configuration.

1. Select **Device** > **Setup** > **Operations**.
2. Click **Reboot Device**.

Session Distribution Policies

Session distribution policies define how PA-5200 and PA-7000 Series firewalls distribute security processing (App-ID, Content-ID, URL filtering, SSL decryption, and IPSec) among dataplane processors (DPs) on the firewall. Each policy is specifically designed for a certain type of network environment and firewall configuration to ensure that the firewall distributes sessions with maximum efficiency. For example, the Hash session distribution policy is best fit for environments that use large scale source NAT.

The number of DPs on a firewall varies based on the firewall model:

Firewall Model	Dataplane Processor(s)
PA-7000 Series	Depends on the number of installed Network Processing Cards (NPCs). Each NPC has multiple dataplane processors (DPs) and you can install multiple NPCs in the firewall.
PA-5220 firewall	1  *The PA-5220 firewall has only one DP so sessions distribution policies do not have an effect. Leave the policy set to the default (round-robin).*
PA-5250 firewall	2
PA-5260 firewall	3

The following topics provide information about the available session distribution policies, how to change an active policy, and how to view session distribution statistics.

- Session Distribution Policy Descriptions
- Change the Session Distribution Policy and View Statistics

Session Distribution Policy Descriptions

The following table provides information about Session Distribution Policies to help you decide which policy best fits your environment and firewall configuration.

© 2018 Palo Alto Networks, Inc.

Session Distribution Policy	Description
Fixed	Allows you to specify the dataplane processor (DP) that the firewall will use for security processing. Use this policy for debugging purposes.
Hash	The firewall distributes sessions based on a hash of the source address or destination address. Hash based distribution improves the efficiency of NAT address resource management and reduces latency for NAT session setup by avoiding potential IP address or port conflicts. Use this policy in environments that use large scale source NAT with dynamic IP translation or Dynamic IP and Port translation or both. When using dynamic IP translation, select the **`source`** address option. When using dynamic IP and port translation, select the **`destination`** address option.
Ingres-slot (default on PA-7000 Series firewalls)	(PA-7000 Series firewalls only) New sessions are assigned to a DP on the same NPC on which the first packet of the session arrived. The selection of the DP is based on the session-load algorithm but, in this case, sessions are limited to the DPs on the ingress NPC. Depending on the traffic and network topology, this policy generally decreases the odds that traffic will need to traverse the switch fabric. Use this policy to reduce latency if both ingress and egress are on the same NPC. If the firewall has a mix of NPCs (PA-7000 20G and PA-7000 20GXM for example), this policy can isolate the increased capacity to the corresponding NPCs and help to isolate the impact of NPC failures.
Random	The firewall randomly selects a DP for session processing.
Round-robin (default on PA-5200 Series firewalls)	The firewall selects the dataplane processor based on a round-robin algorithm between active dataplanes so that input, output, and security processing functions are shared among all dataplanes. Use this policy in low to medium demand environments where a simple and predictable load balancing algorithm will suffice. In high demand environments, we recommend that you use the session-load algorithm.
Session-load	This policy is similar to the round-robin policy but uses a weight-based algorithm to determine how to distribute sessions to achieve balance among the DPs. Because of the variability in the lifetime of a session, the DPs may not always experience an equal load. For example, if the firewall has three DPs and DP0 is at 25% of capacity, DP1 is at 25%, and DP2 is at 50%, new session assignment will be weighted towards the DP with the lower capacities. This helps improve load balancing over time.

© 2018 Palo Alto Networks, Inc.

Session Distribution Policy	Description
	Use this policy in environments where sessions are distributed across multiple NPC slots, such as in an inter-slot aggregate interface group or environments with asymmetric forwarding. You can also use this policy or the ingress-slot policy if the firewall has a combination of NPCs with different session capacities (such as a combination of PA-7000 20G and PA-7000 20GXM NPCs).
Symmetric-hash	(PA-5200 Series and PA-7000 Series firewalls running PAN-OS 8.0 or later) The firewall selects the DP by a hash of sorted source and destination IP addresses. This policy provides the same results for server-to-client (s2c) and client-to-server (c2s) traffic (assuming the firewall does not use NAT). Use this policy in high-demand IPSec or GTP deployments. With these protocols, each direction is treated as a unidirectional flow where the flow tuples cannot be derived from each other. This policy improves performance and reduces latency by ensuring that both directions are assigned to the same DP, which removes the need for inter-DP communication.

Change the Session Distribution Policy and View Statistics

The following table describes how to view and change the active Session Distribution Policies and describes how to view session statistics for each dataplane processor (DP) in the firewall.

Task	Command

Show the active session distribution policy.	Use the **show session distribution policy** command to view the active session distribution policy.
	The following output is from a PA-7080 firewall with four NPCs installed in slots 2, 10, 11, and 12 with the ingress-slot distribution policy enabled:
	`> show session distribution policy`
	`Ownership Distribution Policy: ingress-slot`
	`Flow Enabled Line Cards: [2, 10, 11, 12]Packet Processing Enabled Line Cards: [2, 10, 11, 12]`

Change the active session distribution policy.	Use the **set session distribution-policy *<policy>*** command to change the active session distribution policy. For example, to select the session-load policy, enter the following command: `> set session distribution-policy session-load`

© 2018 Palo Alto Networks, Inc.

Task	Command

View session distribution statistics.	Use the **show session distribution statistics** command to view the dataplane processors (DPs) on the firewall and the number of sessions on each active DP.
	The following output is from a PA-7080 firewall:
	`> show session distribution statistics`
	`DP Active Dispatched Dispatched/sec`
	`----------------------------------------`
	`s1dp0 78698 7829818 1473`
	`s1dp1 78775 7831384 1535`
	`s3dp0 7796 736639 1488`
	`s3dp1 7707 737026 1442`
	The `DP Active column` lists each dataplane on the installed NPCs. The first two characters indicate the slot number and the last three characters indicate the dataplane number. For example, s1dp0 indicates dataplane 0 on the NPC in slot 1 and s1dp1 indicates dataplane 1 on the NPC in slot1.
	The `Dispatched` column shows the total number of sessions that the dataplane processed since the last time the firewall rebooted.
	The `Dispatched/sec` column indicates the dispatch rate. If you add the numbers in the `Dispatched` column, the total equals the number of active sessions on the firewall. You can also view the total number of active sessions by running the **show session info CLI** command.
	The PA-5200 Series firewall output will look similar, except that the number of DPs depends on the model and there is only one NPC slot (s1).

Prevent TCP Split Handshake Session Establishment

You can configure a TCP Split Handshake Drop in a Zone Protection profile to prevent TCP sessions from being established unless they use the standard three-way handshake. This task assumes that you assigned a security zone for the interface where you want to prevent TCP split handshakes from establishing a session.

STEP 1 | Configure a Zone Protection profile to prevent TCP sessions that use anything other than a three-way handshake to establish a session.

1. Select **Network** > **Network Profiles** > **Zone Protection** and **Add** a new profile (or select an existing profile).
2. If creating a new profile, enter a **Name** for the profile and an optional **Description**.
3. Select **Packet Based Attack Protection** > **TCP Drop** and select **Split Handshake**.
4. Click **OK**.

STEP 2 | Apply the profile to one or more security zones.

1. Select **Network** > **Zones** and select the zone where you want to assign the zone protection profile.
2. In the Zone window, from the **Zone Protection Profile** drop-down, select the profile you configured in the previous step.

© 2018 Palo Alto Networks, Inc.

Alternatively, you could start creating a new profile here by clicking **Zone Protection Profile**, in which case you would continue accordingly.

3. Click **OK**.
4. (Optional) Repeat steps 1-3 to apply the profile to additional zones.

STEP 3 | Commit your changes.

Click **OK** and **Commit**.

© 2018 Palo Alto Networks, Inc.

Tunnel Content Inspection

The firewall can inspect the traffic content of cleartext tunnel protocols:

- Generic Routing Encapsulation (GRE) (RFC 2784)
- Non-encrypted IPSec traffic [NULL Encryption Algorithm for IPSec (RFC 2410) and transport mode AH IPSec]
- General Packet Radio Service (GPRS) Tunneling Protocol for User Data (GTP-U)

You can use tunnel content inspection to enforce Security, DoS Protection, and QoS policies on traffic in these types of tunnels and traffic nested within another cleartext tunnel (for example, a Null Encrypted IPSec tunnel inside a GRE tunnel). You can view tunnel inspection logs and tunnel activity in the ACC to verify that tunneled traffic complies with your corporate security and usage policies.

All firewall models support tunnel content inspection of GRE and non-encrypted IPSec. Tunnel content inspection of GTP-U is supported only on the PA-5200 Series and VM-Series firewalls. The firewalls don't terminate GRE, non-encrypted IPSec, or GTP-U tunnels.

Tunnel content inspection is for cleartext tunnels, not for VPN or LSVPN tunnels, which carry encrypted traffic.

- Tunnel Content Inspection Overview
- Configure Tunnel Content Inspection
- View Inspected Tunnel Activity
- View Tunnel Information in Logs
- Create a Custom Report Based on Tagged Tunnel Traffic

Tunnel Content Inspection Overview

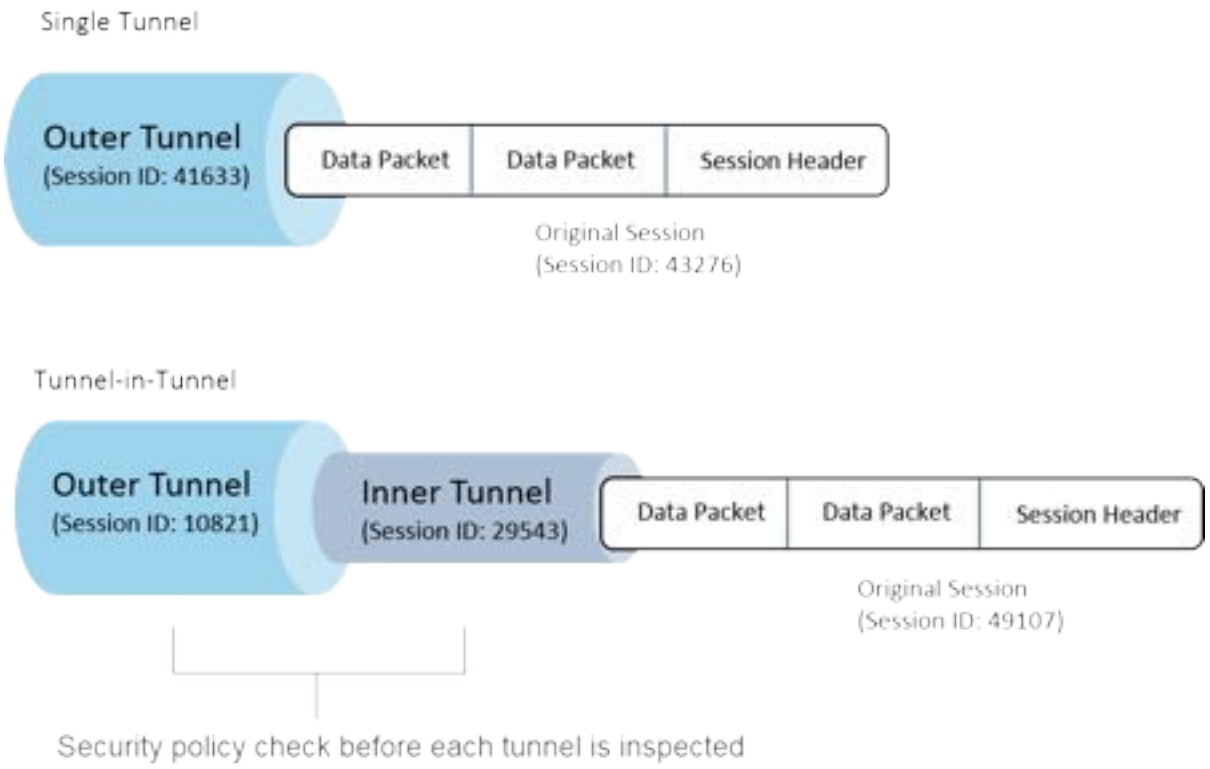
Your firewall can inspect tunnel content anywhere on the network where you do not have the opportunity to terminate the tunnel first. As long as the firewall is in the path of a GRE, non-encrypted IPSec, or GTP-U tunnel, the firewall can inspect the tunnel content.

- Enterprise customers who want tunnel content inspection can have some or all of the traffic on the firewall tunneled using GRE or non-encrypted IPSec. For security, QoS, and reporting reasons, you want to inspect the traffic inside the tunnel.
- Service Provider customers use GTP-U to tunnel data traffic from mobile devices. You want to inspect the inner content without terminating the tunnel protocol, and you want to record user data from your users.

The firewall supports tunnel content inspection on Ethernet interfaces and subinterfaces, AE interfaces, VLAN interfaces, and VPN and LSVPN tunnel interfaces. (The cleartext tunnel that the firewall inspects can be inside a VPN or LSVPN tunnel that terminates at the firewall, hence a VPN or LSVPN tunnel interface. In other words, when the firewall is a VPN or LSVPN endpoint, the firewall can inspect the traffic of any non-encrypted tunnel protocols that tunnel content inspection supports.)

Tunnel content inspection is supported in Layer 3, Layer 2, virtual wire, and tap deployments. Tunnel content inspection works on shared gateways and on virtual system-to-virtual system communications.

© 2018 Palo Alto Networks, Inc.



The preceding figure illustrates the two levels of tunnel inspection the firewall can perform. When a firewall configured with Tunnel Inspection policy rules receives a packet:

- The firewall first performs a Security policy check to determine whether the tunnel protocol (Application) in the packet is permitted or denied. (IPv4 and IPv6 packets are supported protocols inside the tunnel.)
- If the Security policy allows the packet, the firewall matches the packet to a Tunnel Inspection policy rule based on source zone, source address, source user, destination zone, and destination address. The Tunnel Inspection policy rule determines the tunnel protocols that the firewall inspects, the maximum level of encapsulation allowed (a single tunnel or a tunnel within a tunnel), whether to allow packets containing a tunnel protocol that doesn't pass strict header inspection per RFC 2780, and whether to allow packets containing unknown protocols.
- If the packet passes the Tunnel Inspection policy rule's match criteria, the firewall inspects the inner content, which is subject to your Security policy (required) and optional policies you can specify. (The supported policy types for the original session are listed in the following table).
- If the firewall instead finds another tunnel, the firewall recursively parses the packet for the second header and is now at level two of encapsulation, so the second tunnel inspection policy rule, which matches a tunnel zone, must allow a maximum tunnel inspection level of two levels for the firewall to continue processing the packet.
    - If your rule allows two levels of inspection, the firewall performs a Security policy check on this inner tunnel and then the Tunnel Inspection policy check. The tunnel protocol you use in an inner tunnel can differ from the tunnel protocol you use in the outer tunnel.
    - If your rule doesn't allow two levels of inspection, the firewall bases its action on whether you configured it to drop packets that have more levels of encapsulation than the maximum tunnel inspection level you configured.

By default, the content encapsulated in a tunnel belongs to the same security zone as the tunnel, and is subject to the Security policy rules that protect that zone. However, you can configure a *tunnel zone*, which gives you the flexibility to configure Security policy rules for inside content that differ from the Security

© 2018 Palo Alto Networks, Inc.

policy rules for the tunnel. If you use a different tunnel inspection policy for the tunnel zone, it must always have a maximum tunnel inspection level of two levels because by definition the firewall is looking at the second level of encapsulation.



Although tunnel content inspection works on shared gateways and on virtual system-to-virtual system communications, you can't assign tunnel zones to shared gateways or virtual system-to-virtual system communications; they are subject to the same Security policy rules as the zones to which they belong.

The following table indicates with a check mark which types of policy you can apply to an outer tunnel session, an inner tunnel session, and the inside, original session:

Policy Type	Outer Tunnel Session	Inner Tunnel Session	Inside, Original Session
App-Override	–	–	✓
DoS Protection	✓	✓	✓
NAT	✓	–	–
Policy-Based Forwarding (PBF) and Symmetric Return	✓	–	–
QoS	–	–	✓
Security (required)	✓	✓	✓
User-ID	✓	✓	✓
Zone Protection	✓	✓	✓

The inner tunnel sessions and outer tunnel sessions count toward the maximum session capacity for the firewall model.

When you enable or edit a Tunnel Inspection policy (to add a protocol, increase maximum levels of inspection, or enable security options), you affect existing tunnel sessions. The firewall treats existing TCP sessions inside the tunnel as non-SYN TCP flows. To prevent the firewall from dropping all existing sessions when you enable or edit a Tunnel Inspection policy, you can create a Zone Protection profile that disables **Reject Non-SYN TCP** and apply the profile to the zones that control your tunnel's security policies. The task to Configure Tunnel Content Inspection includes these steps.

The firewall doesn't support a Tunnel Inspection policy rule that matches traffic for a tunnel that terminates on the firewall; the firewall discards packets that match the inner tunnel session. For example, when an IPSec tunnel terminates on the firewall, don't create a Tunnel Inspection policy rule that matches the tunnel you terminate. The firewall already inspects the inner tunnel traffic, so no Tunnel Inspection policy rule is necessary.

You can View Inspected Tunnel Activity on the ACC or View Tunnel Information in Logs. To facilitate quick viewing, configure a Monitor tag so you can monitor tunnel activity and filter log results by that tag.

The ACC tunnel activity provides data in various views. For the Tunnel ID Usage, Tunnel Monitor Tag, and Tunnel Application Usage, the data for **bytes**, **sessions**, **threats**, **content**, and **URLs** come from the Traffic Summary database. For the Tunnel User, Tunneled Source IP and Tunneled Destination IP Activity, data for

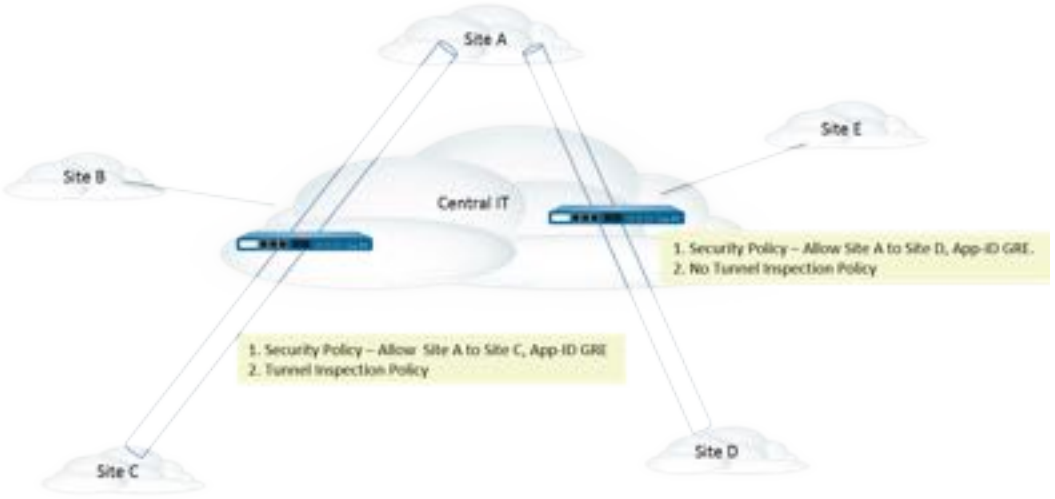
© 2018 Palo Alto Networks, Inc.

bytes and **sessions** come from Traffic Summary database, data for **threats** come from the Threat Summary, data for **URLs** come from the URL Summary, and data for **contents** come from the Data database, which is a subset of the Threat logs.

If you enable NetFlow on the interface, NetFlow will capture statistics for the outer tunnel only, to avoid double-counting (counting bytes of both outer and inner flows).

For the Tunnel Inspection policy rule and tunnel zone capacities for your firewall model, see the Product Selection tool.

The following figure illustrates a corporation that runs multiple divisions and uses different Security policies and a Tunnel Inspection policy. A Central IT team provides connectivity between regions. A tunnel connects Site A to Site C; another tunnel connects Site A to Site D. Central IT places a firewall in the path of each tunnel; the firewall in the tunnel between Sites A and C performs tunnel inspection; the firewall in the tunnel between Sites A and D has no tunnel inspection policy because the traffic is very sensitive.



Configure Tunnel Content Inspection

Perform this task to configure tunnel content inspection for a tunnel protocol that you allow in a tunnel.

STEP 1 | Create a Security policy to allow packets through the tunnel from the source zone to the destination zone that use a specific application, such as the GRE application.

Create a Security Policy Rule

The firewall can create tunnel inspection logs at the start or end of a session. When you specify Actions for the Security policy rule, select Log at Session Start for long-lived tunnel sessions such as GRE sessions.

STEP 2 | Create a Tunnel Inspection policy rule.

1. Select **Policies** > **Tunnel Inspection** and **Add** a policy rule.
2. On the **General** tab, enter a Tunnel Inspection policy rule **Name**, beginning with an alphanumeric character and containing zero or more alphanumeric, underscore (_), hyphen (-), dot (.), and space characters.
3. (Optional) Enter a **Description**.
4. (Optional) Specify a **Tag** that identifies the packets that are subject to the Tunnel Inspection policy rule, for reporting and logging purposes.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Specify the criteria that determine the source of packets to which the Tunnel Inspection policy rule applies.

1. Select the **Source** tab.
2. **Add** a **Source Zone** from the list of zones. The default is **Any** zone.
3. (Optional) **Add** a **Source Address**. You can enter an IPv4 or IPv6 address, an address group, or a Geo Region address object. The default is **Any** source address.
4. (Optional) Select **Negate** to choose any addresses except the specified ones.
5. (Optional) **Add** a **Source User**. The default is **any** source user. **Known-user** is a user who has authenticated; an **Unknown** user has not authenticated.

STEP 4 | Specify the criteria that determine the destination of packets to which the Tunnel Inspection policy rule applies.

1. Select the **Destination** tab.
2. **Add** a **Destination Zone** from the list of zones. The default is **Any** zone.
3. (Optional) **Add** a **Destination Address**. You can enter an IPv4 or IPv6 address, an address group, or a Geo Region address object. The default is **Any** destination address.

   You can also configure a new Address or Address Group.
4. (Optional) Select **Negate** to choose any addresses except the specified ones.

STEP 5 | Specify the tunnel protocols the firewall will inspect for this rule.

1. Select the **Inspection** tab.
2. **Add** one or more tunnel **Protocols** that you want the firewall to inspect:
   - **GRE**—Firewall inspects packets that use Generic Route Encapsulation in the tunnel.
   - **GTP-U**—Firewall inspects packets that use General Packet Radio Service (GPRS) Tunneling Protocol for User Data (GTP-U) in the tunnel.
   - **Non-encrypted IPSec**—Firewall inspects packets that use non-encrypted IPSec (Null Encrypted IPSec or transport mode AH IPSec) in the tunnel.

STEP 6 | Specify how many levels of encapsulation the firewall inspects and the conditions under which the firewall drops a packet.

1. Select **Inspect Options**.
2. Select the **Maximum Tunnel Inspection Levels** that the firewall will inspect:
   - **One Level**—Firewall inspects content that is in the outer tunnel only (default).
   - **Two Levels (Tunnel In Tunnel)**—Firewall inspects content that is in the outer tunnel and content that is in the inner tunnel.
3. Select the following to specify whether the firewall drops a packet under this condition:
   1. **Drop packet if over maximum tunnel inspection level**—Firewall drops a packet that contains more levels of encapsulation than are configured for **Maximum Tunnel Inspection Levels**.
   2. **Drop packet if tunnel protocol fails strict header check**—Firewall drops a packet that contains a tunnel protocol that uses a header that is non-compliant with the RFC for the protocol. Non-compliant headers can indicate suspicious packets. This option causes the firewall to verify GRE headers against RFC 2890.

      If your firewall is tunneling GRE with a device that implements a version of GRE older than RFC 2890, you shouldn't enable the option to Drop packet if tunnel protocol fails strict header check.
   3. **Drop packet if unknown protocol inside tunnel**—Firewall drops a packet that contains a protocol inside the tunnel that the firewall can't identify.

© 2018 Palo Alto Networks, Inc.

For example, if this option is selected, the firewall drops encrypted IPSec packets that match the Tunnel Inspection policy rule because the firewall can't read them. Thus, you can allow IPSec packets, and the firewall will allow only null-encrypted IPSec and AH IPSec packets.

4. Click **OK**.

STEP 7 | Manage Tunnel Inspection policy rules.

Use the following to manage Tunnel Inspection policy rules:

- (Filter field)—Displays only the tunnel policy rules named in the filter field.
- **Delete**—Removes selected tunnel policy rules.
- **Clone**—An alternative to the **Add** button; duplicates the selected rule with a new name, which you can then revise.
- **Enable**—Enables the selected tunnel policy rules.
- **Disable**—Disables the selected tunnel policy rules.
- **Move**—Moves the selected tunnel policy rules up or down in the list; packets are evaluated against the rules in order from the top down.
- **Highlight Unused Rules**—Highlights tunnel policy rules that no packets have matched since the last time the firewall was restarted.

STEP 8 | (Optional) Create a tunnel source zone and tunnel destination zone for tunnel content and configure a Security policy rule for each zone.



The best practice is to create tunnel zones for your tunnel traffic. Thus, the firewall creates separate sessions for tunneled and non-tunneled packets that have the same five-tuple (source IP address and port, destination IP address and port, and protocol).



Assigning tunnel zones to tunnel traffic on a PA-5200 Series firewall causes the firewall to do tunnel inspection in software; tunnel inspection isn't offloaded to hardware.

1. If you want tunnel content to be subject to different Security policy rules from the Security policy rules for the zone of the outer tunnel (configured earlier), select **Network** > **Zones** and **Add** a **Name** for the Tunnel Source Zone.
2. For **Location**, select the virtual system.
3. For **Type**, select **Tunnel**.
4. Click **OK**.
5. Repeat these substeps to create the Tunnel Destination Zone.
6. Configure a Security Policy Rule for the Tunnel Source Zone.



Because you might not know the originator of the tunnel traffic or the direction of the traffic flow and you don't want to inadvertently prohibit traffic for an application through the tunnel, specify both tunnel zones as the Source Zone and specify both tunnel zones as the Destination Zone in your Security policy rule, or select Any for both the source and destination zones; then specify the Applications.

7. Configure a Security Policy Rule for the Tunnel Destination Zone. The tip for configuring a Security policy rule for the Tunnel Source Zone applies to the Tunnel Destination Zone also.

STEP 9 | (Optional) Specify the Tunnel Source Zone and Tunnel Destination Zone for the inner content.

1. Specify the Tunnel Source Zone and Tunnel Destination Zone you just added as the zones for the inner content. Select **Policies** > **Tunnel Inspection** and on the **General** tab, select the **Name** of the Tunnel Inspection policy rule you created.
2. Select **Inspection**.
3. Select **Security Options**.

© 2018 Palo Alto Networks, Inc.

4. Select **Enable Security Options** to cause the inner content source to belong to the **Tunnel Source Zone** you specify, and to cause the inner content destination to belong to the **Tunnel Destination Zone** you specify. (Default is disabled.)

   If you don't **Enable Security Options**, the inner content source belongs to the same source zone as the outer tunnel source, and the inner content destination belongs to the same destination zone as the outer tunnel destination, and they are therefore subject to the same Security policy rules that apply to those outer zones.
5. For **Tunnel Source Zone**, select the appropriate tunnel zone you created in the prior step so that the policies associated with that zone apply to the tunnel source zone. Otherwise, by default the inner content will use the same source zone that is used in the outer tunnel, and the policies of the outer tunnel source zone apply to the inner content source zone also.
6. For **Tunnel Destination Zone**, select the appropriate tunnel zone you created in the prior step so that the policies associated with that zone apply to the tunnel destination zone. Otherwise, by default the inner content will use the same destination zone that is used in the outer tunnel, and the policies of the outer tunnel destination zone apply to the inner content destination zone also.

   If you configure a Tunnel Source Zone and Tunnel Destination Zone for the tunnel inspection policy rule, you should configure a specific Source Zone (in Step 3) and a specific Destination Zone (in Step 4) in the match criteria of the tunnel inspection policy rule, instead of specifying a Source Zone of Any and a Destination Zone of Any. This tip ensures the direction of zone reassignment corresponds to the parent zones.
7. Click **OK**.

STEP 10 | (Optional) If you enabled **Rematch Sessions** (**Device** > **Setup** > **Session**), ensure the firewall doesn't drop existing sessions when you create or revise a Tunnel Inspection policy, by disabling **Reject Non-SYN TCP** for the zones that control your tunnel's Security policies.

The firewall displays the following warning when you:

- Create a Tunnel Inspection policy rule.
- Edit a Tunnel Inspection policy rule by adding a **Protocol** or by increasing the **Maximum Tunnel Inspection Levels** from **One Level** to **Two Levels**.
- **Enable Security Options** in the **Security Options** tab by either adding new zones or changing one zone to another zone.



Warning: Enabling tunnel inspection policies on existing tunnel sessions will cause existing TCP sessions inside the tunnel to be treated as non-syn-tcp flows. To ensure existing sessions are not dropped when the tunnel inspection policy is enabled, set the Reject Non-SYN TCP setting for the zone(s) to no using a Zone Protection profile and apply it to the zones that control the tunnel's security policies. Once the existing sessions have been recognized by the firewall, you can re-enable the Reject Non-SYN TCP setting by setting it to yes or global.

1. Select **Network** > **Network Profiles** > **Zone Protection** and **Add** a profile.
2. Enter a **Name** for the profile.
3. Select **Packet Based Attack Protection** > **TCP Drop**.
4. For **Reject Non-SYN TCP**, select **no**.
5. Click **OK**.
6. Select **Network** > **Zones** and select the zone that controls your tunnel's security policies.
7. For **Zone Protection Profile**, select the Zone Protection profile you just created.
8. Click **OK**.
9. Repeat the prior three steps in this section to apply the Zone Protection profile to additional zones that control your tunnel's Security policies.

© 2018 Palo Alto Networks, Inc.

10. After the firewall has recognized the existing sessions, you can re-enable **Reject Non-SYN TCP** by setting it to **yes** or **global**.

STEP 11 | Tag tunnel traffic for aggregated logging and reporting across firewalls or outside the firewall.



If you tag tunnel traffic, you can later filter on the Monitor Tag in the Tunnel Inspection log and use the ACC to view tunnel activity based on Monitor Tag.

1. Select **Policies** > **Tunnel Inspection** and select the Tunnel Inspection policy rule you created.
2. Select **Inspection** > **Monitor Options**.
3. Enter a **Monitor Name** to group similar traffic together for purposes of logging and reporting.
4. Enter a **Monitor Tag (number)** to group similar traffic together for logging and reporting (range is 1-16,777,215). The tag number is globally defined.
5. Click **OK**.

STEP 12 | (Optional) Limit fragmentation of traffic in a tunnel.

1. Select **Network** > **Network Profiles** > **Zone Protection** and **Add** a profile by **Name**.
2. Enter a **Description**.
3. Select **Packet Based Attack Protection** > **IP Drop** > **Fragmented traffic**.
4. Click **OK**.
5. Select **Network** > **Zones** and select the tunnel zone where you want to limit fragmentation.
6. For **Zone Protection Profile**, select the profile you just created to apply the Zone Protection profile to the tunnel zone.
7. Click **OK**.

STEP 13 | Commit your changes.

Click **Commit**.

View Inspected Tunnel Activity

Perform the following task to view activity of inspected tunnels.

STEP 1 | Select **ACC** and select a **Virtual System** or **All** virtual systems.

STEP 2 | Select Tunnel Activity.

STEP 3 | Select a Time period to view, such as Last 24 Hrs or Last 30 Days.

STEP 4 | For Global Filters, click the + or - buttons to use ACC Filters on tunnel activity.

STEP 5 | View inspected tunnel activity; you can display and sort data in each window by **bytes**, **sessions**, **threats**, **content**, or **URLs**. Each window displays a different aspect of tunnel data in graph and table format:

- **Tunnel ID Usage**—Each tunnel protocol lists the Tunnel IDs of tunnels using that protocol. Tables provide totals of Bytes, Sessions, Threats, Content, and URLs for the protocol. Hover over the tunnel ID to get a breakdown per tunnel ID.
- **Tunnel Monitor Tag**—Each tunnel protocol lists tunnel monitor tags of tunnels using that tag. Tables provide totals of Bytes, Sessions, Threats, Content, and URLs for the tag and for the protocol. Hover over the tunnel monitor tag to get a breakdown per tag.

© 2018 Palo Alto Networks, Inc.

- **Tunneled Application Usage**—Application categories graphically display types of applications grouped into media, general interest, collaboration, and networking, and color-coded by their risk. The Application tables also include a count of users per application.
- **Tunneled User Activity**—Displays a graph of bytes sent and bytes received, for example, along an x-axis of date and time. Hover over a point on the graph to view data at that point. The Source User and Destination User table provides data per user.
- **Tunneled Source IP Activity**—Displays graphs and tables of bytes, sessions, and threats, for example, from an Attacker at an IP address. Hover over a point on the graph to view data at that point.
- **Tunneled Destination IP Activity**—Displays graphs and tables based on destination IP addresses. View threats per Victim at an IP address, for example. Hover over a point on the graph to view data at that point.

View Tunnel Information in Logs

You can view Tunnel Inspection logs themselves or view tunnel inspection information in other types of logs.

- View Tunnel inspection logs.
  1. Select **Monitor** > **Logs** > **Tunnel Inspection** and view the log data, noting the tunnel **Applications** used in your traffic and any high counts for packets failing Strict Checking of headers, for example.
  2. Click the Detailed Log View icon  to see details about a log.

- View other logs for tunnel inspection information.
  1. Select **Monitor** > **Logs**.
  2. Select **Traffic**, **Threat**, **URL Filtering**, **WildFire Submissions**, **Data Filtering**, or **Unified**.
  3. For a log entry, click the Detailed Log View icon  .
  4. In the Flags window, see if the **Tunnel Inspected** flag is checked. A Tunnel Inspected flag indicates the firewall used a Tunnel Inspection policy rule to inspect the inside content or inner tunnel. Parent Session information refers to an outer tunnel (relative to an inner tunnel) or an inner tunnel (relative to inside content).

     On the **Traffic**, **Threat**, **URL Filtering**, **WildFire Submissions**, **Data Filtering** logs, only direct parent information appears in the Detailed Log View of the inner session log, no tunnel log information. If you configured two levels of tunnel inspection, you can select the parent session of this direct parent to view the second parent log. (You must monitor the **Tunnel Inspection** log as shown in the prior step to view tunnel log information.)
  5. If you are viewing the log for an inside session that is tunnel inspected, click the **View Parent Session** link in the General section to see the outside session information.

Create a Custom Report Based on Tagged Tunnel Traffic

You can create a report to gather information based on the tag you applied to tunnel traffic.

STEP 1 | Select **Monitor** > **Manage Custom Reports** and click **Add**.

STEP 2 | For Database, select the Traffic, Threat, URL, Data Filtering, or WildFire Submissions log.

STEP 3 | For Available Columns, select Flags and Monitor Tag, along with other data you want in the report.

© 2018 Palo Alto Networks, Inc.

See Generate Custom Reports for details about creating a custom report.

Policy

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Authentication, Denial of Service (DoS), and Zone protection policies. All these different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. The following topics describe how to work with policy:

> Policy Types
> Security Policy
> Policy Objects
> Security Profiles
> Enumeration of Rules Within a Rulebase
> Move or Clone a Policy Rule or Object to a Different Virtual System
> Use Tags to Group and Visually Distinguish Objects
> Use an External Dynamic List in Policy
> Register IP Addresses and Tags Dynamically
> Monitor Changes in the Virtual Environment
> CLI Commands for Dynamic IP Addresses and Tags
> Identify Users Connected through a Proxy Server
> Policy-Based Forwarding

© 2018 Palo Alto Networks, Inc.

Policy Types

The Palo Alto Networks next-generation firewall supports a variety of policy types that work together to safely enable applications on your network.

Policy Type	Description
Security	Determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service. For more details, see Security Policy.
NAT	Instruct the firewall which packets need translation and how to do the translation. The firewall supports both source address and/or port translation and destination address and/or port translation. For more details, see NAT.
QoS	Identify traffic requiring QoS treatment (either preferential treatment or bandwidth-limiting) using a defined parameter or multiple parameters and assign it a class. For more details, see Quality of Service.
Policy Based Forwarding	Identify traffic that should use a different egress interface than the one that would normally be used based on the routing table. For details, see Policy-Based Forwarding.
Decryption	Identify encrypted traffic that you want to inspect for visibility, control, and granular security. For more details, see Decryption.
Application Override	Identify sessions that you do not want processed by the App-ID engine, which is a Layer-7 inspection. Traffic matching an application override policy forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4. For more details, see Manage Custom or Unknown Applications.
Authentication	Identify traffic that requires users to authenticate. For more details, see Authentication Policy.
DoS Protection	Identify potential denial-of-service (DoS) attacks and take protective action in response to rule matches. DoS Protection Profiles.

© 2018 Palo Alto Networks, Inc.

Security Policy

Security policy protects network assets from threats and disruptions and aids in optimally allocating network resources for enhancing productivity and efficiency in business processes. On the Palo Alto Networks firewall, individual Security policy rules determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service.



To ensure that end users authenticate when they try to access your network resources, the firewall evaluates Authentication Policy *before Security policy.*

All traffic passing through the firewall is matched against a session and each session is matched against a Security policy rule. When a session match occurs, the firewall applies the matching Security policy rule to bi-directional traffic (client to server and server to client) in that session. For traffic that doesn't match any defined rules, the default rules apply. The default rules—displayed at the bottom of the security rulebase—are predefined to allow all intrazone (within the zone) traffic and deny all interzone (between zones) traffic. Although these rules are part of the pre-defined configuration and are read-only by default, you can override them and change a limited number of settings, including the tags, action (allow or block), log settings, and security profiles.

Security policy rules are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log, if logging is enabled for that rule. The logging options are configurable for each rule, and can for example be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.

- Components of a Security Policy Rule
- Security Policy Actions
- Create a Security Policy Rule

Components of a Security Policy Rule

The Security policy rule construct permits a combination of the required and optional fields as detailed in the following table:

Required/ Optional	Field	Description
Required	**Name**	A label that supports up to 31 characters, used to identify the rule.
	Rule Type	Specifies whether the rule applies to traffic within a zone, between zones, or both: • **universal** (default)—Applies the rule to all matching interzone and intrazone traffic in the specified source and destination zones. For example, if you create a universal rule with source zones A and B and destination zones A and B, the rule would apply to all traffic within zone A, all traffic within zone B, and all traffic from zone A to zone B and all traffic from zone B to zone A. • **intrazone**—Applies the rule to all matching traffic within the specified source zones (you cannot specify a destination zone for intrazone rules). For example, if you set the source zone to A and B, the rule

© 2018 Palo Alto Networks, Inc.

Required/ Optional	Field	Description
		would apply to all traffic within zone A and all traffic within zone B, but not to traffic between zones A and B. • **interzone**—Applies the rule to all matching traffic between the specified source and destination zones. For example, if you set the source zone to A, B, and C and the destination zone to A and B, the rule would apply to traffic from zone A to zone B, from zone B to zone A, from zone C to zone A, and from zone C to zone B, but not traffic within zones A, B, or C.
	Source Zone	The zone from which the traffic originates.
	Destination Zone	The zone at which the traffic terminates. If you use NAT, make sure to always reference the post-NAT zone.
	Application	The application which you wish to control. The firewall uses App-ID, the traffic classification technology, to identify traffic on your network. App-ID provides application control and visibility in creating security policies that block unknown applications, while enabling, inspecting, and shaping those that are allowed.
	Action	Specifies an *Allow* or *Block* action for the traffic based on the criteria you define in the rule. When you configure the firewall to block traffic, it either resets the connection or silently drops packets. To provide a better user experience, you can configure granular options to block traffic instead of silently dropping packets, which can cause some applications to break and appear unresponsive to the user. For more details, see Security Policy Actions.
Optional	**Tag**	A keyword or phrase that allows you to filter security rules. This is handy when you have defined many rules and wish to then review those that are tagged with a keyword such as *IT-sanctioned applications* or *High-risk applications*.
	Description	A text field, up to 255 characters, used to describe the rule.
	Source IP Address	Define host IP or FQDN, subnet, named groups, or country-based enforcement. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address).
	Destination IP Address	The location or destination for the traffic. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address).
	User	The user or group of users for whom the policy applies. You must have User-ID enabled on the zone. To enable User-ID, see User-ID Overview.
	URL Category	Using the URL Category as match criteria allows you to customize security profiles (Antivirus, Anti-Spyware, Vulnerability, File-Blocking, Data Filtering, and DoS) on a per-URL-category basis. For example, you can prevent.exe file download/upload for URL categories that represent higher risk while allowing them for other categories. This functionality

© 2018 Palo Alto Networks, Inc.

Required/ Optional	Field	Description
		also allows you to attach schedules to specific URL categories (allow social-media websites during lunch & after-hours), mark certain URL categories with QoS (financial, medical, and business), and select different log forwarding profiles on a per-URL-category-basis. Although you can manually configure URL categories on your firewall, to take advantage of the dynamic URL categorization updates available on the Palo Alto Networks firewalls, you must purchase a URL filtering license.  *To block or allow traffic based on URL category, you must apply a URL Filtering profile to the security policy rules. Define the URL Category as Any and attach a URL Filtering profile to the security policy. See* Set Up a Basic Security Policy *for information on using the default profiles in your security policy and see* Control Access to Web Content *for more details.*
	Service	Allows you to select a Layer 4 (TCP or UDP) port for the application. You can choose *any*, specify a port, or use *application-default* to permit use of the standards-based port for the application. For example, for applications with well- known port numbers such as DNS, the *application-default* option will match against DNS traffic only on TCP port 53. You can also add a custom application and define the ports that the application can use.  *For inbound allow rules (for example, from untrust to trust), using application-default prevents applications from running on unusual ports and protocols. Application-default is the default option; while the firewall still checks for all applications on all ports, with this configuration, applications are only allowed on their standard ports/ protocols.*
	Security Profiles	Provide additional protection from threats, vulnerabilities, and data leaks. Security profiles are only evaluated for rules that have an *allow* action.
	HIP Profile (for GlobalProtect)	Allows you to identify clients with Host Information Profile (HIP) and then enforce access privileges.
	Options	Allow you to define logging for the session, log forwarding settings, change Quality of Service (QoS) markings for packets that match the rule, and schedule when (day and time) the security rule should be in effect.

Security Policy Actions

For traffic that matches the attributes defined in a security policy, you can apply the following actions:

© 2018 Palo Alto Networks, Inc.

Action	Description
Allow (default)	Allows the traffic.
Deny	Blocks traffic and enforces the default *Deny Action* defined for the application that is being denied. To view the deny action defined by default for an application, view the application details in **Objects** > **Applications** or check the application details in Applipedia.
Drop	Silently drops the traffic; for an application, it overrides the default deny action. A TCP reset is not sent to the host/application. For Layer 3 interfaces, to optionally send an ICMP unreachable response to the client, set Action: **Drop** and enable the **Send ICMP Unreachable** check box. When enabled, the firewall sends the ICMP code for *communication with the destination is administratively prohibited*—ICMPv4: Type 3, Code 13; ICMPv6: Type 1, Code 1.
Reset client	Sends a TCP reset to the client-side device.
Reset server	Sends a TCP reset to the server-side device.
Reset both	Sends a TCP reset to both the client-side and server-side devices.

A reset is sent only after a session is formed. If the session is blocked before a 3-way handshake is completed, the firewall will not send the reset.

For a TCP session with a reset action, the firewall does not send an ICMP Unreachable response.

For a UDP session with a drop or reset action, if the **ICMP Unreachable** check box is selected, the firewall sends an ICMP message to the client.

Create a Security Policy Rule

STEP 1 | (Optional) Delete the default Security policy rule.

By default, the firewall includes a security rule named *rule1* that allows all traffic from Trust zone to Untrust zone. You can either delete the rule or modify the rule to reflect your zone naming conventions.

STEP 2 | Add a rule.

1. Select **Policies** > **Security** and click **Add**.
2. Enter a descriptive **Name** for the rule in the **General** tab.
3. Select a **Rule Type**.

STEP 3 | Define the matching criteria for the source fields in the packet.

1. In the **Source** tab, select a **Source Zone**.
2. Specify a **Source IP Address** or leave the value set to **any**.
3. Specify a Source **User** or leave the value set to **any**.

STEP 4 | Define the matching criteria for the destination fields in the packet.

1. In the **Destination** tab, set the **Destination Zone**.
2. Specify a **Destination IP Address** or leave the value set to **any**.

© 2018 Palo Alto Networks, Inc.



As a best practice, consider using address objects in the Destination Address field to enable access to specific servers or groups of servers only, particularly for services such as DNS and SMTP that are commonly exploited. By restricting users to specific destination server addresses you can prevent data exfiltration and command and control traffic from establishing communication through techniques such as DNS tunneling.

STEP 5 | Specify the application the rule will allow or block.



As a best practice, always use application-based security policy rules instead of port based rules and always set the Service to application-default unless you are using a more restrictive list of ports than the standard ports for an application.

1. In the **Applications** tab, **Add** the **Application** to safely enable. You can select multiple applications, or use application groups or application filters.
2. In the **Service/URL Category** tab, keep the Service set to **application-default** to ensure that any applications the rule allows are only allowed on their standard ports.

STEP 6 | (Optional) Specify a URL category as match criteria for the rule.

In the **Service/URL Category** tab, select the **URL Category**.

If you select a URL category, only web traffic will match the rule and only if the traffic is to the specified category.

STEP 7 | Define what action you want the firewall to take for traffic that matches the rule.

In the **Actions** tab, select an **Action**. See Security Policy Actions for a description of each action.

STEP 8 | Configure the log settings.

- By default, the rule is set to **Log at Session End**. You can clear this setting if you don't want any logs generated when traffic matches this rule, or select **Log at Session Start** for more detailed logging.
- Select a **Log Forwarding** profile.



As a best practice, do not select the check box to Disable Server Response Inspection (DSRI). Selecting this option prevents the firewall from inspecting packets from the server to the client. For the best security posture, the firewall must inspect both the client-to-server flows and the server-to-client flows to detect and prevent threats.

STEP 9 | Attach security profiles to enable the firewall to scan all allowed traffic for threats.



See Create Best Practice Security Profiles to learn how to create security profiles that protect your network from both known and unknown threats.

In the **Actions** tab, select **Profiles** from the **Profile Type** drop-down and then select the individual security profiles to attach to the rule.

Alternatively, select **Group** from the **Profile Type** drop-down and select a security **Group Profile** to attach.

STEP 10 | Save the policy rule to the running configuration on the firewall.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

STEP 11 | To verify that you have set up your basic policies effectively, test whether your security policy rules are being evaluated and determine which security policy rule applies to a traffic flow.

To verify the policy rule that matches a flow, use the following CLI command:

```
test security-policy-match source <IP_address> destination <IP_address>
 destination port <port_number> protocol <protocol_number>
```

The output displays the best rule that matches the source and destination IP address specified in the CLI command.

For example, to verify the policy rule that will be applied for a server in the data center with the IP address 208.90.56.11 when it accesses the Microsoft update server:

```
test security-policy-match source 208.80.56.11 destination 176.9.45.70
 destination-port 80 protocol 6
"Updates-DC to Internet" {
        from data_center_applications;
        source any;
        source-region any;
        to untrust;
        destination any;
        destination-region any;
        user any;
        category any;
application/service[dns/tcp/any/53 dns/udp/any/53 dns/udp/any/5353 ms-
update/tcp/any/80 ms-update/tcp/any/443];
action allow;
        terminal yes;
```


© 2018 Palo Alto Networks, Inc.

Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With policy objects that are a collective unit, you can reference the object in security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

You can create the following policy objects on the firewall:

Policy Object	Description
Address/Address Group, Region	Allow you to group specific source or destination addresses that require the same policy enforcement. The address object can include an IPv4 or IPv6 address (single IP, range, subnet) or the FQDN. Alternatively, a region can be defined by the latitude and longitude coordinates or you can select a country and define an IP address or IP range. You can then group a collection of address objects to create an address group object. You can also use dynamic address groups to dynamically update IP addresses in environments where host IP addresses change frequently.
User/User Group	Allow you to create a list of users from the local database or an external database and group them.
Application Group and Application Filter	An Application Filter allows you to filter applications dynamically. It allows you to filter, and save a group of applications using the attributes defined in the application database on the firewall. For example, you can Create an Application Filter on page 518 by one or more attributes—category, sub-category, technology, risk, characteristics. With an application filter, when a content update occurs, any new applications that match your filter criteria are automatically added to your saved application filter. An Application Group allows you to create a static group of specific applications that you want to group together for a group of users or for a particular service, or to achieve a particular policy goal. See Create an Application Group on page 518.
Service/Service Groups	Allows you to specify the source and destination ports and protocol that a service can use. The firewall includes two pre-defined services—service-http and service-https— that use TCP ports 80 and 8080 for HTTP, and TCP port 443 for HTTPS. You can however, create any custom service on any TCP/UDP port of your choice to restrict application usage to specific ports on your network (in other words, you can define the default port for the application).  *To view the standard ports used by an application, in Objects > Applications search for the application and click the link. A succinct description displays.*

© 2018 Palo Alto Networks, Inc.

Security Profiles

While 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 allowed applications for threats, such as viruses, malware, spyware, and DDOS attacks. When traffic matches the allow rule defined in the security policy, the security profile(s) that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.



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. See Set Up a Basic Security Policy for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, see Create Best Practice Security Profiles for the Internet Gateway to learn how you can create custom profiles and get best practice recommendations.

You can add 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 security policies in one step (or included in security policies by default, if you choose to set up a default security profile group).

Profile Type	Description
Antivirus Profiles	Antivirus profiles protect against viruses, worms, and trojans as well as spyware downloads. Using a stream-based malware prevention engine, which inspects traffic the moment the first packet is received, the Palo Alto Networks antivirus solution can provide protection for clients without significantly impacting the performance of the firewall. This profile scans for a wide variety of malware in executables, PDF files, HTML and JavaScript viruses, including support for scanning inside compressed files and data encoding schemes. If you have enabled Decryption on the firewall, the profile also enables scanning of decrypted content. The default profile inspects all of the listed protocol decoders for viruses, and generates alerts for SMTP, IMAP, and POP3 protocols while blocking for FTP, HTTP, and 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**—Permits the application traffic. • **Alert**—Generates an alert for each application traffic flow. The alert is saved in the threat log. • **Drop**—Drops the application traffic. • **Reset Client**—For TCP, resets the client-side connection. For UDP, drops the connection. • **Reset Server**—For TCP, resets the server-side connection. For UDP, drops the connection. • **Reset Both**—For TCP, resets the connection on both client and server ends. For UDP, drops the connection.

© 2018 Palo Alto Networks, Inc.

Profile Type	Description
	Customized profiles can be used to minimize antivirus inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the internet, as well as the traffic sent to highly sensitive destinations, such as server farms. 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 by Threat Prevention subscribers on a daily basis (sub-hourly for WildFire subscribers).
Anti-Spyware Profiles	Anti-Spyware profiles blocks spyware on compromised hosts from trying to phone-home or beacon out to external command-and-control (C2) servers, allowing you to detect 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 traffic received from an untrusted zone, such as internet-facing zones. 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**—Uses the default action for every signature, as specified by Palo Alto Networks when the signature is created. • **Strict**—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. When the firewall detects a threat event, you can configure the following actions in an Anti-Spyware profile: • **Default**—For each threat signature and Anti-Spyware 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**—Permits the application traffic • **Alert**—Generates an alert for each application traffic flow. The alert is saved in the threat log. • **Drop**—Drops the application traffic. • **Reset Client**—For TCP, resets the client-side connection. For UDP, drops the connection. • **Reset Server**—For TCP, resets the server-side connection. For UDP, drops the connection. • **Reset Both**—For TCP, resets the connection on both client and server ends. For UDP, drops the connection. • **Block IP**— This action blocks traffic from either a source or a source-destination pair. It is configurable for a specified period of time. In addition, you can enable the DNS Sinkholing action in Anti-Spyware profiles to enable the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP

© 2018 Palo Alto Networks, Inc.

Profile Type	Description
	address that you define. This feature helps to identify infected hosts on the protected network 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 are most likely infected with malware. Anti-Spyware and Vulnerability Protection profiles are configured similarly.
Vulnerability Protection Profiles	Vulnerability Protection profiles stop attempts to exploit system flaws or gain unauthorized access to systems. While 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 can also create exceptions, which allow you to change the response to a specific signature. To configure how the firewall responds to a threat, see Anti-Spyware Profiles for a list of supported actions.
URL Filtering Profiles	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 profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories.
Data Filtering Profiles	Data filtering profiles prevent sensitive information such as credit card or social security numbers from leaving a protected network. The data filtering profile also allows you to filter on key words, such as a sensitive project name or the word confidential. It is important to focus your profile on the desired file types to reduce false positives. For example, you may only want to search Word documents or Excel spreadsheets. You may also only want to scan web-browsing traffic, or FTP. You can create custom data pattern objects and attach them to a Data Filtering profile to define the type of information on which you want to filter. Create data pattern objects based on: • **Predefined Patterns**—Filter for credit card and social security numbers (with or without dashes) using predefined patterns. • **Regular Expressions**—Filter for a string of characters. • **File Properties**—Filter for file properties and values based on file type.  *If you're using a third-party, endpoint data loss prevention (DLP) solutions to populate file properties to indicate sensitive content, this option enables the firewall to enforce your DLP policy.* To get started, Set Up Data Filtering.

© 2018 Palo Alto Networks, Inc.

Profile Type	Description
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 and/or download and you can specify which applications will be subject to the file blocking profile. You can also configure custom block pages that will appear when a user attempts to download the specified file type. This allows the user to take a moment to consider whether or not they want to 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 quickly enable best practice file blocking settings: • **basic file blocking**—Attach this profile to the Security policy rules that allow traffic to and from less sensitive applications to block files that are commonly included in malware attack campaigns or that have no real use case for upload/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, .bat. Additionally, it prompts users to acknowledge when they attempt to download encrypted-rar or encrypted-zip files. This rule alerts on all other file types to give you complete visibility into all file types coming in and out of your network. • **strict file blocking**—Use this stricter profile on the Security policy rules that allow access to your most sensitive applications. This profile blocks the same file types as the other profile, and additionally blocks flash, .tar, multi-level encoding, .cab, .msi, encrypted-rar, and encrypted-zip files. Configure a file blocking profile with the following actions: • **Alert**—When the specified file type is detected, a log is generated in the data filtering log. • **Block**—When the specified file type is detected, the file is blocked and a customizable block page is presented to the user. A log is also generated in the data filtering log. • **Continue**—When 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. Because this type of forwarding action requires user interaction, it is only applicable for web traffic. To get started, Set Up File Blocking.
WildFire Analysis Profiles	Use a WildFire analysis profile to enable the firewall to forward unknown files or email links for WildFireanalysis. Specify files to be forwarded for analysis based on application, file type, and transmission direction (upload or download). Files or email links matched to the profile rule are forwarded 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 existing antivirus signatures, in addition to unknown files. You can also use the WildFire analysis profiles to set up a Wildfire hybrid cloud deployment. If you are using a WildFire appliance to analyze sensitive

© 2018 Palo Alto Networks, Inc.

Profile Type	Description
	files locally (such as PDFs), you can specify for less sensitive files 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 the WildFire appliance and the WildFire cloud for analysis allows you to benefit from a prompt verdict for files that have already been processed by the cloud, and for files that are not supported for appliance analysis, and frees up the appliance capacity to process sensitive content.
DoS Protection Profiles	DoS protection profiles provide detailed control for Denial of Service (DoS) protection policies. DoS policies allow you to control the number of sessions between interfaces, zones, addresses, and countries based on aggregate sessions or source and/or destination IP addresses. There are two DoS protection mechanisms that the Palo Alto Networks firewalls support. • **Flood Protection**—Detects and prevents attacks where the network is flooded with packets resulting in too many half-open sessions and/or services being unable to respond to each request. In this case the source address of the attack is usually spoofed. See DoS Protection Against Flooding of New Sessions. • **Resource Protection**— Detects and prevent session exhaustion attacks. In this type of attack, a large number of hosts (bots) are used to establish as many fully established sessions as possible to consume all of a system's resources. You can enable both types of protection mechanisms in a single DoS protection profile. The DoS profile is used to specify the type of action to take and details on matching criteria for the DoS policy. The DoS profile defines settings for SYN, UDP, and ICMP floods, can enable resource protect and defines the maximum number of concurrent connections. After you configure the DoS protection profile, you then attach it to a DoS policy. When configuring DoS protection, it is important to analyze your environment in order to set the correct thresholds and due to some of the complexities of defining DoS protection policies, this guide will not go into detailed examples.
Zone Protection Profiles	Zone Protection Profiles provide additional protection between specific network zones in order to protect the zones against attack. The profile must be applied to the entire zone, so it is important to carefully test the profiles in order to prevent issues that may arise with the normal traffic traversing the zones. When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session.
Security Profile Group	A security profile group is a set of security profiles that can be treated as a unit and then easily added to security policies. Profiles that are often assigned together can be added to profile groups to simplify the creation of security policies. You can also setup a default security profile group—new security policies will use the settings defined in the default profile group to check and control traffic that matches the security policy. Name a security profile group default to allow the profiles in that group to be added to new security policies by default. This allows you to consistently include your organization's

© 2018 Palo Alto Networks, Inc.

Profile Type	Description
	preferred profile settings in new policies automatically, without having to manually add security profiles each time you create new rules. See Create a Security Profile Group and Set Up or Override a Default Security Profile Group *For recommendations on the best-practice settings for security profiles, see Best Practice Security Profiles.*

Create a Security Profile Group

Use the following steps to create a security profile group and add it to a security policy.

STEP 1 | Create a security profile group.

If you name the group `default`, the firewall will automatically attach it to any new rules you create. This is a time saver if you have a preferred set of security profiles that you want to make sure get attached to every new rule.

1. Select **Objects** > **Security Profile Groups** and **Add** a new security profile group.
2. Give the profile group a descriptive **Name**, for example, Threats.
3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.
4. Add existing profiles to the group.



5. Click **OK** to save the profile group.

STEP 2 | Add a security profile group to a security policy.

1. Select **Policies** > **Security** and **Add** or modify a security policy rule.
2. Select the **Actions** tab.
3. In the Profile Setting section, select **Group** for the **Profile Type**.
4. In the **Group Profile** drop-down, select the group you created (for example, select the best-practice group):



5. Click **OK** to save the policy and **Commit** your changes.

STEP 3 | Save your changes.

© 2018 Palo Alto Networks, Inc.

Click **Commit**.

Set Up or Override a Default Security Profile Group

Use the following options to set up a default security profile group to be used in new security policies, or to override an existing default group. When an administrator creates a new security policy, the default profile group will be automatically selected as the policy's profile settings, and traffic matching the policy will be checked according to the settings defined in the profile group (the administrator can choose to manually select different profile settings if desired). Use the following options to set up a default security profile group or to override your default settings.

If no default security profile exists, the profile settings for a new security policy are set to None by default.

- Create a security profile group.
  1. Select **Objects** > **Security Profile Groups** and Add a new security profile group.
  2. Give the profile group a descriptive **Name**, for example, Threats.
  3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.
  4. Add existing profiles to the group. For details on creating profiles, see Security Profiles.



  5. Click **OK** to save the profile group.
  6. Add the security profile group to a security policy.
  7. **Add** or modify a security policy rule and select the **Actions** tab.
  8. Select **Group** for the **Profile Type**.
  9. In the **Group Profile** drop-down, select the group you created (for example, select the Threats group):



  10. Click **OK** to save the policy and **Commit** your changes.

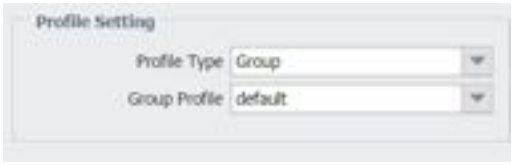
- Set up a default security profile group.
  1. Select **Objects** > **Security Profile Groups** and add a new security profile group or modify an existing security profile group.
  2. **Name** the security profile group **`default`**:



  3. Click **OK** and **Commit**.
  4. Confirm that the default security profile group is included in new security policies by default:
     1. Select **Policies** > **Security** and **Add** a new security policy.

© 2018 Palo Alto Networks, Inc.

2. Select the **Actions** tab and view the **Profile Setting** fields:



   By default, the new security policy correctly shows the **Profile Type** set to Group and the default **Group Profile** is selected.

- Override a default security profile group.

  If you have an existing default security profile group, and you do not want that set of profiles to be attached to a new security policy, you can continue to modify the Profile Setting fields according to your preference. Begin by selecting a different Profile Type for your policy (**Policies** > **Security** > **Security Policy Rule** > **Actions**).

© 2018 Palo Alto Networks, Inc.

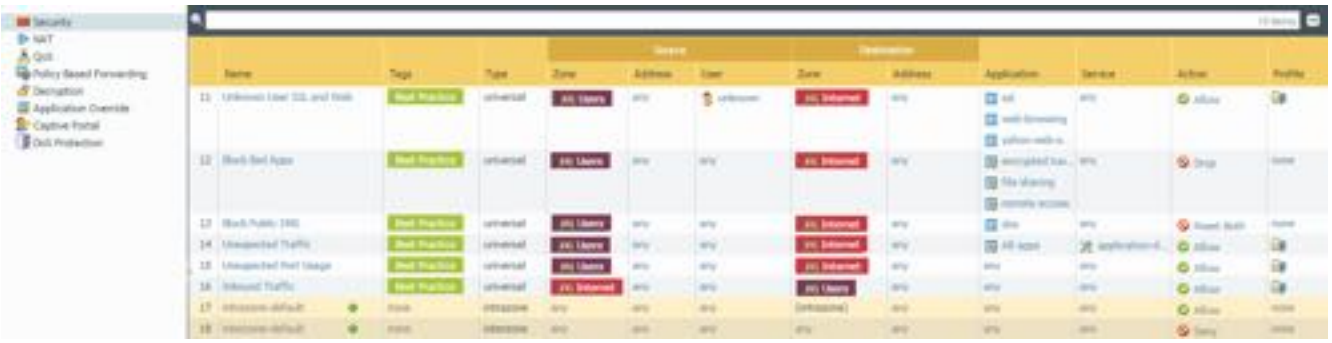
Enumeration of Rules Within a Rulebase

Each rule within a rulebase is automatically numbered and the ordering adjusts as rules are moved or reordered. When filtering rules to find rules that match the specified filter(s), each rule is listed with its number in the context of the complete set of rules in the rulebase and its place in the evaluation order.

On Panorama, pre-rules, post-rules, and default rules are independently numbered. When Panorama pushes rules to a firewall, the rule numbering reflects the hierarchy and evaluation order of shared rules, device group pre-rules, firewall rules, device group post-rules, and default rules. The **Preview Rules** option in Panorama displays an ordered list view of the total number of rules on a firewall.

- View the numbered list of rules on the firewall.

  Select **Policies** and any rulebase under it. For example, **Policies** > **Security**. The left-most column in the table displays the rule number.



- View the numbered list of rules on Panorama.

  Select **Policies** and any rulebase under it. For example, **Policies** > **Security** > **Pre-rules**.



- After you push the rules from Panorama, view the complete list of rules with numbers on the firewall.

  From the web interface of the firewall, select **Policies** and pick any rulebase under it. For example, select **Policies** > **Security** and view the complete set of numbered rules that the firewall will evaluate.

© 2018 Palo Alto Networks, Inc.



© 2018 Palo Alto Networks, Inc.

Move or Clone a Policy Rule or Object to a Different Virtual System

On a firewall that has more than one virtual system (vsys), you can move or clone policy rules and objects to a different vsys or to the Shared location. Moving and cloning save you the effort of deleting, recreating, or renaming rules and objects. If the policy rule or object that you will move or clone from a vsys has references to objects in that vsys, move or clone the referenced objects also. If the references are to shared objects, you do not have to include those when moving or cloning. You can Use Global Find to Search the Firewall or Panorama Management Server for references.

When cloning multiple policy rules, the order by which you select the rules will determine the order they are copied to the device group. For example, if you have rules 1-4 and your selection order is 2-1-4-3, the device group where these rules will be cloned will display the rules in the same order you selected. However, you can reorganize the rules as you see fit once they have been successfully copied.

STEP 1 | Select the policy type (for example, **Policy** > **Security**) or object type (for example, **Objects** > **Addresses**).

STEP 2 | Select the **Virtual System** and select one or more policy rules or objects.

STEP 3 | Perform one of the following steps:

- Select **Move** > **Move to other vsys** (for policy rules).
- Click **Move** (for objects).
- Click **Clone** (for policy rules or objects).

STEP 4 | In the **Destination** drop-down, select the new virtual system or **Shared**.

STEP 5 | (Policy rules only) Select the **Rule order**:

- **Move top** (default)—The rule will come before all other rules.
- **Move bottom**—The rule will come after all other rules.
- **Before rule**—In the adjacent drop-down, select the rule that comes after the Selected Rules.
- **After rule**—In the adjacent drop-down, select the rule that comes before the Selected Rules.

STEP 6 | The **Error out on first detected error in validation** check box is selected by default. The firewall stops performing the checks for the move or clone action when it finds the first error, and displays just this error. For example, if an error occurs when the **Destination** vsys doesn't have an object that the policy rule you are moving references, the firewall will display the error and stop any further validation. When you move or clone multiple items at once, selecting this check box will allow you to find one error at a time and troubleshoot it. If you clear the check box, the firewall collects and displays a list of errors. If there are any errors in validation, the object is not moved or cloned until you fix all the errors.

STEP 7 | Click **OK** to start the error validation. If the firewall displays errors, fix them and retry the move or clone operation. If the firewall doesn't find errors, the object is moved or cloned successfully. After the operation finishes, click **Commit**.

© 2018 Palo Alto Networks, Inc.

Use Tags to Group and Visually Distinguish Objects

You can tag objects to group related items and add color to the tag in order to visually distinguish them for easy scanning. You can create tags for the following objects: address objects, address groups, zones, service groups, and policy rules.

The firewall and Panorama support both static tags and dynamic tags. Dynamic tags are registered from a variety of sources and are not displayed with the static tags because dynamic tags are not part of the firewall/Panorama configuration. See Register IP Addresses and Tags Dynamically on page 1021 for information on registering tags dynamically. The tags discussed in this section are statically added and are part of the configuration.

You can apply one or more tags to objects and to policy rules, up to a maximum of 64 tags per object. Panorama supports a maximum of 10,000 tags, which you can apportion across Panorama (shared and device groups) and the managed firewalls (including firewalls with multiple virtual systems).

- Create and Apply Tags on page 1006
- Modify Tags on page 1007
- Use the Tag Browser on page 1007

Create and Apply Tags

STEP 1 | Create tags.

To tag a zone, you must create a tag with the same name as the zone. When the zone is attached in policy rules, the tag color automatically displays as the background color against the zone name.

1. Select **Objects** > **Tags**.
2. On Panorama or a multiple virtual system firewall, select the **Device Group** or the **Virtual System** to to make the tag available.
3. Click **Add** and enter a **Name** to identify the tag, or select a zone name from the drop-down to create a tag for a zone. The maximum length is 127 characters.
4. (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.
5. (Optional) Assign one of the 17 predefined colors to the tag. By default, **Color** is **None**.

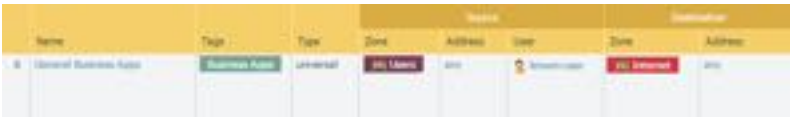


6. Click **OK** and **Commit** to save the changes.

STEP 2 | Apply tags to policy.

1. Select **Policies** and any rulebase under it.
2. Click **Add** to create a policy rule and use the tagged objects you created in Step 1.
3. Verify that the tags are in use.

© 2018 Palo Alto Networks, Inc.



STEP 3 | Apply tags to an address object, address group, service, or service group.

1. Create the object.

   For example to create a service group, select **Objects** > **Service Groups** > **Add**.
2. Select a tag from the **Tags** drop-down or enter a name in the field to create a new tag.

   To edit a tag or add color to the tag, see Modify Tags.

Modify Tags

- Select **Objects** > **Tags** to perform any of the following operations with tags:
  - Click the link in the **Name** column to edit the properties of a tag.
  - Select a tag in the table, and click **Delete** to remove the tag from the firewall.
  - Click **Clone** to create a duplicate tag with the same properties. A numerical suffix is added to the tag name. For example, FTP-1.

  For details on creating tags, see Create and Apply Tags on page 1006. For information on working with tags, see Use the Tag Browser on page 1007.

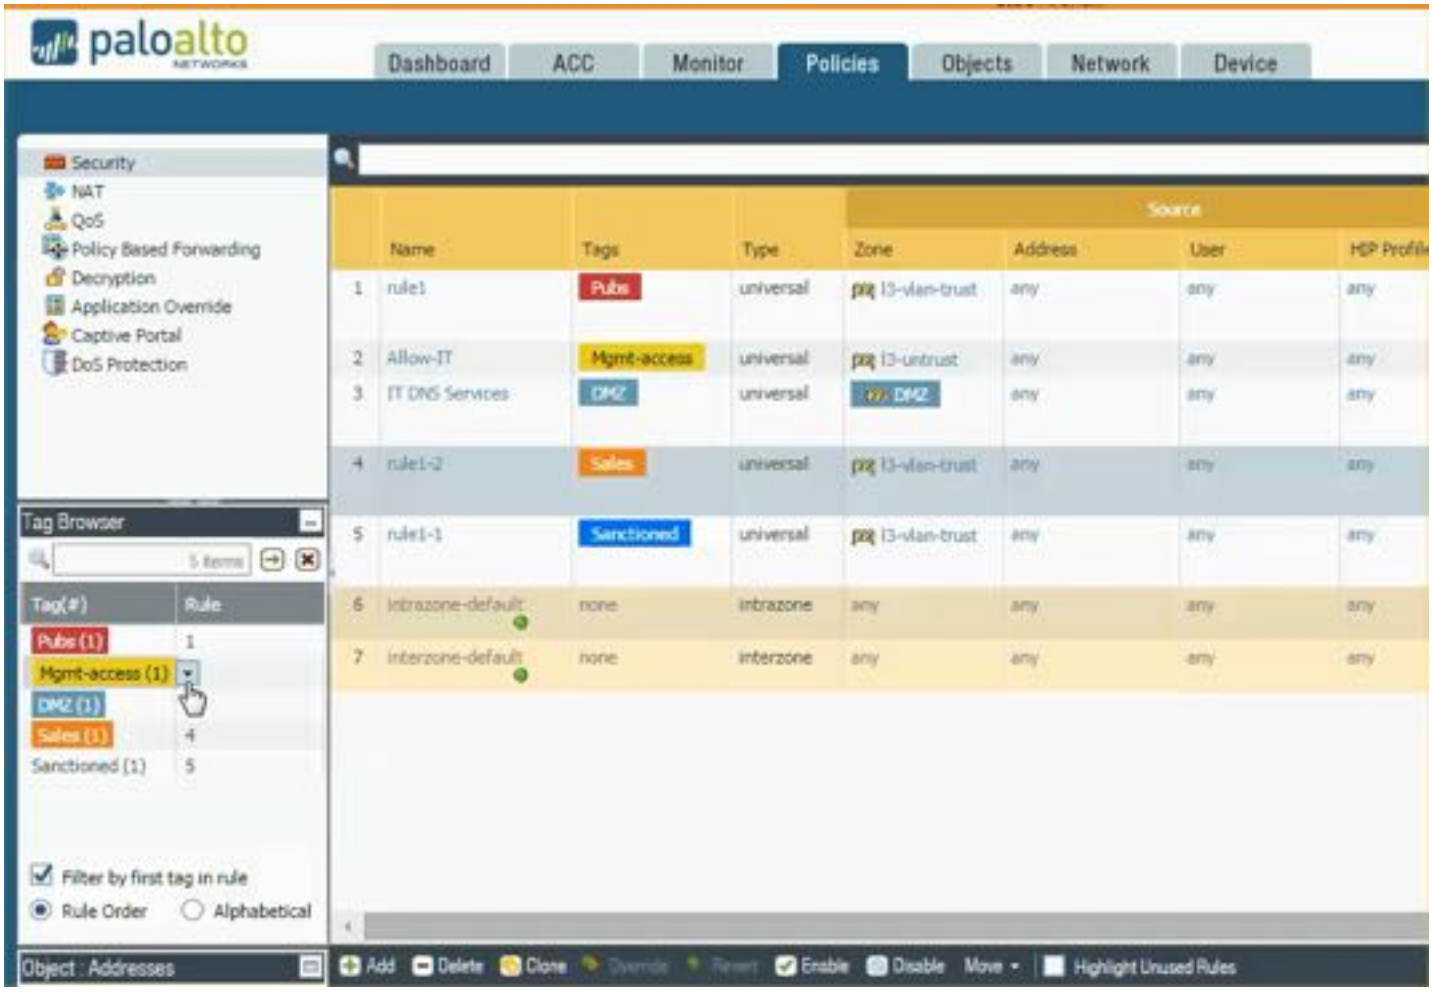
Use the Tag Browser

The tag browser provides a way to view all the tags used within a rulebase. In rulebases with a large number of rules, the tag browser simplifies the display by presenting the tags, the color code, and the rule numbers in which the tags are used.

It also allows you to group rules using the first tag applied to the rule. As a best practice, use the first tag to identify the primary purpose for a rule. For example, the first tag can identify a rule by a high-level function such as best practice, or internet access or IT sanctioned applications or high-risk applications. In the tag browser, when you **Filter by first tag in rule**, you can easily identify gaps in coverage and move rules or add new rules within the rulebase. All the changes are saved to the candidate configuration until you commit the changes on the firewall and make them a part of the running configuration.

For firewalls that are managed by Panorama, the tags applied to pre-rules and post-rules that have been pushed from Panorama, display in a green background and are demarcated with green lines so that you can identify these tags from the local tags on the firewall.

© 2018 Palo Alto Networks, Inc.



- Explore the tag browser.
  1. Access the **Tag Browser** on the left pane of the **Policies** tab. The tag browser displays the tags that have been used in the rules for the selected rulebase, for example **Policies** > **Security**.
  2. **Tag (#)**—Displays the label and the rule number or range of numbers in which the tag is used contiguously. Hover over the label to see the location where the rule was defined, it can be inherited from a shared location, a device group, or a virtual system.
  3. **Rule**—Lists the rule number or range of numbers associated with the tags.
  4. Sort the tags.
     - **Filter by first tag in rule**—Sorts rules using the first tag applied to each rule in the rulebase. This view is particularly useful if you want to narrow the list and view related rules that might be spread around the rulebase. For example if the first tag in each rule denotes its function—best practices, administration, web-access, data center access, proxy—you can narrow the result and scan the rules based on function.
     - **Rule Order**—Sorts the tags in the order of appearance within the selected rulebase. When displayed in order of appearance, tags used in contiguous rules are grouped. The rule number with which the tag is associated is displayed along with the tag name.
     - **Alphabetical**—Sorts the tags in alphabetical order within the selected rulebase. The display lists the tag name and color (if a color is assigned) and the number of times it is used within the rulebase.

© 2018 Palo Alto Networks, Inc.

The label **None** represents rules without any tags; it does not display rule numbers for untagged rules. When you select **None**, the right pane is filtered to display rules that have no tags assigned to them.

5. **Clear**—Clears the filter on the currently selected tags in the search bar.



6. **Search bar**—To search for a tag, enter the term and click the green arrow icon to apply the filter. It also displays the total number of tags in the rulebase and the number of selected tags.
7. Expand or collapse the tag browser.



- Tag a rule.
  1. Select a rule on the right pane.
  2. Do one of the following:
     - Select a tag in the tag browser and select **Apply the Tag to the Selection(s)** from the drop-down.



     - Drag and drop tag(s) from the tag browser on to the Tags column of the rule. When you drop a tag, a confirmation dialog displays.
  3. **Commit** the changes.

- View rules that match the selected tags.

  You can filter rules based on tags with an AND or an OR operator.

  - OR filter: To view rules that have specific tags, select one or more tags in the tag browser; the right pane only displays the rules that include any of the currently selected tags.
  - AND filter: To view rules that have all the selected tags, hover over the number associated with the tag in the **Rule** column of the tag browser and select **Filter**. Repeat to add more tags.

  Click the apply filter icon in the search bar on the right pane. The results are displayed using an AND operator.

© 2018 Palo Alto Networks, Inc.



- View the currently selected tags.

  To view the currently selected tags, hover over the **Clear** label in the tag browser.



- Untag a rule.

  Hover over the rule number associated with a tag in the **Rule** column of the tag browser and select **Untag Rule(s)**. Confirm that you want to remove the selected tag from the rule. **Commit** the changes.

- Reorder rules using tags.

  Select one or more tags and hover over the rule number in the Rule column of the tag browser and select **Move Rule(s)**.

  Select a tag from the drop-down in the move rule window and select whether you want to **Move Before** or **Move After** the tag selected in the drop-down. **Commit** the changes.

- Add a new rule that applies the selected tags.

  Select one or more tags and hover over the rule number in the **Rule** column of the tag browser, and select **Add New Rule**. Define the rule and **Commit** the changes.

  The numerical order of the new rule varies by whether you selected a rule on the right pane. If you did not select a rule on the right pane, the new rule will be added after the rule to which the selected tag(s) belongs. Otherwise, the new rule is added after the selected rule.

- Search for a tag.

  In the tag browser, enter the first few letters of the tag name you want to search for and click the Apply Filter icon. The tags that match your input will display.

© 2018 Palo Alto Networks, Inc.

Use an External Dynamic List in Policy

An external dynamic list (formerly called dynamic block list) is a text file that you or another source hosts on an external web server so that the firewall can import objects—IP addresses, URLs, domains—to enforce policy on the entries in the list. As the list is updated, the firewall dynamically imports the list at the configured interval and enforces policy without the need to make a configuration change or a commit on the firewall.

- External Dynamic List on page 1011
- Formatting Guidelines for an External Dynamic List on page 1012
- Palo Alto Networks Malicious IP Address Feeds on page 1013
- Configure the Firewall to Access an External Dynamic List on page 1014
- Retrieve an External Dynamic List from the Web Server on page 1016
- View External Dynamic List Entries on page 1016
- Exclude Entries from an External Dynamic List on page 1017
- Enforce Policy on an External Dynamic List on page 1017
- Find External Dynamic Lists That Failed Authentication on page 1019
- Disable Authentication for an External Dynamic List on page 1020

External Dynamic List

An External Dynamic List is a text file that is hosted on an external web server so that the firewall can import objects—IP addresses, URLs, domains—included in the list and enforce policy. To enforce policy on the entries included in the external dynamic list, you must reference the list in a supported policy rule or profile. As you modify the list, the firewall dynamically imports the list at the configured interval and enforces policy without the need to make a configuration change or a commit on the firewall. If the web server is unreachable, the firewall will use the last successfully retrieved list for enforcing policy until the connection is restored with the web server, but only if the list is not secured with SSL. To retrieve the external dynamic list, the firewall uses interface configured with the **Palo Alto Networks Services** service route.

The firewall supports four types of external dynamic lists:

- **IP Address**—The firewall typically enforces policy for a source or destination IP address that is defined as a static object on the firewall (see Use an External Dynamic List of Type IP or Predefined IP as a Source or Destination Address Object in a Security Policy Rule.) If you need agility in enforcing policy for a list of source or destination IP addresses that emerge ad hoc, you can use an external dynamic list of type IP address as a source or destination address object in policy rules, and configure the firewall to deny or allow access to the IP addresses (IPv4 and IPv6 address, IP range and IP subnets) included in the list. The firewall treats an external dynamic list of type IP address as an address object; all the IP addresses included in a list are handled as one address object.
- **Predefined IP Address**—A predefined IP address list is a type of IP address list that refers to any of the two Palo Alto Networks Malicious IP Address Feeds that have fixed or "predefined" contents. These feeds are automatically added to your firewall if you have an active Threat Prevention license. A predefined IP address list can also refer to any external dynamic list you create that uses a Palo Alto Networks IP address feed as a source.
- **URL**—An external dynamic list of type URL gives you the agility to protect your network from new sources of threat or malware. The firewall handles an external dynamic list with URLs like a custom URL category and you can use this list in two ways:
  - As a match criteria in Security policy rules, Decryption policy rules, and QoS policy rules to allow, deny, decrypt, not decrypt, or allocate bandwidth for the URLs in the custom category.

© 2018 Palo Alto Networks, Inc.

- In a URL Filtering profile where you can define more granular actions, such as continue, alert, or override, before you attach the profile to a Security policy rule (see Use an External Dynamic List in a URL Filtering Profile).

- **Domain**—An external dynamic list of type domain allows you to import custom domain names into the firewall to enforce policy using an Anti-Spyware profile. This capability is very useful if you subscribe to third-party threat intelligence feeds and want to protect your network from new sources of threat or malware as soon as you learn of a malicious domain. For each domain you include in the external dynamic list, the firewall creates a custom DNS-based spyware signature so that you can enable DNS sinkholing. The DNS-based spyware signature is of type spyware with medium severity and each signature is named **`Custom Malicious DNS Query <domain name>`**. For details, see Configure DNS Sinkholing for a List of Custom Domains.

On each firewall model, you can use a maximum of 30 external dynamic lists with unique sources to enforce policy; predefined IP address feeds do not count toward this limit. The external dynamic list limit is not applicable to Panorama. When using Panorama to manage a firewall that is enabled for multiple virtual systems, if you exceed the limit for the firewall, a commit error displays on Panorama. A source is a URL that includes the IP address or hostname, the path, and the filename for the external dynamic list. The firewall matches the URL (complete string) to determine whether a source is unique.

While the firewall does not impose a limit on the number of lists of a specific type, the following limits are enforced:

- IP address—The PA-5000 Series, PA-5200 Series, and the PA-7000 Series firewalls support a maximum of 150,000 total IP addresses; all other models support a maximum of 50,000 total IP addresses. No limits are enforced for the number of IP addresses per list. When the maximum supported IP address limit is reached on the firewall, the firewall generates a syslog message. The IP addresses in predefined IP address lists do not count toward the limit.
- URL and domain—A maximum of 50,000 URLs and 50,000 domains are supported on each model, with no limits enforced on the number of entries per list.

List entries only count toward the firewall limits if they belong to an external dynamic list that is referenced in policy.



- *When parsing the list, the firewall skips entries that do not match the list type, and ignores entries that exceed the maximum number supported for the model. To ensure that the entries do not exceed the limit, check the number of entries currently used in policy. Select Objects > External Dynamic Lists and click List Capacities.*
- *An external dynamic list should not be empty. If you want to stop using the list, remove the reference from the policy rule or profile rather than leave the list empty. If the the list is empty the firewall fails to refresh the list and continues to use the last information it retrieved.*

Formatting Guidelines for an External Dynamic List

An external dynamic list of one type —IP address, URL or Domain—must include entries of that type only. The entries in a predefined IP address list comply with the formatting guidelines for IP address lists.

- IP Address List on page 1012
- Domain List on page 1013
- URL List on page 1013

IP Address List

The external dynamic list can include individual IP addresses, subnet addresses (address/mask), or range of IP addresses. In addition, the block list can include comments and special characters such as **`*`** , **`:`** , **`;`** , **`#`**, or **`/`**. The syntax for each line in the list is **`[IP address, IP/Mask, or IP start range-IP end range] [space] [comment]`**.

© 2018 Palo Alto Networks, Inc.

Enter each IP address/range/subnet in a new line; URLs or domains are not supported in this list. A subnet or an IP address range, such as 92.168.20.0/24 or 192.168.20.40-192.168.20.50, count as one IP address entry and not as multiple IP addresses. If you add comments, the comment must be on the same line as the IP address/range/subnet. The space at the end of the IP address is the delimiter that separates a comment from the IP address.

An example IP address list:

```
192.168.20.10/32
 2001:db8:123:1::1 #test IPv6 address
 192.168.20.0/24 ; test internal subnet
 2001:db8:123:1::/64 test internal IPv6 range
 192.168.20.40-192.168.20.50
```

For an IP address that is blocked, you can display a notification page only if the protocol is HTTP.

Domain List

Enter each domain name in a new line; URLs or IP addresses are not supported in this list. Do not prefix the domain name with the protocol, http:// or https://. Wildcards are not supported.

An example list of domains:

```
www.example.com
 baddomain.com
 qqq.abcedfg.au
```

URL List

See URL Category Exception Lists on page 649.

Palo Alto Networks Malicious IP Address Feeds

With an active Threat Prevention license, Palo Alto Networks provides two feeds with malicious IP addresses that you can use to secure your network against malicious hosts.

- **Palo Alto Networks Known Malicious IP Addresses**—Contains IP addresses that are verified malicious based on WildFire analysis, Unit 42 research, and data gathered from telemetry (Share Threat Intelligence with Palo Alto Networks). Attackers use these IP addresses almost exclusively to distribute malware, initiate command-and-control activity, and launch attacks.
- **Palo Alto Networks High-Risk IP Addresses**—Contains malicious IP addresses from threat advisories issued by trusted third-party organizations. Palo Alto Networks compiles the list of threat advisories, but does not have direct evidence of the maliciousness of the IP addresses.

The firewall receives updates for these feeds through daily antivirus content updates, allowing you to enforce security policy on the firewall based on the latest threat intelligence from Palo Alto Networks. The Palo Alto Networks IP address feeds are predefined, which means that you cannot modify their contents. You can use them as-is (see Enforce Policy on an External Dynamic List), or create a custom external dynamic list that uses either feed as a source (see Configure the Firewall to Access an External Dynamic List) and exclude entries from the list as needed.

© 2018 Palo Alto Networks, Inc.

Configure the Firewall to Access an External Dynamic List

You must establish the connection between the firewall and the source that hosts the external dynamic list before you can Enforce Policy on an External Dynamic List.

STEP 1 | (Optional) Customize the service route that the firewall uses to retrieve external dynamic lists.

Select **Device** > **Setup** > **Services** > **Service Route Configuration** > **Customize** and modify the **External Dynamic Lists** service route.



The firewall does not use the External Dynamic Lists service route to retrieve the Palo Alto Networks Malicious IP Address Feeds; it dynamically receives updates to these feeds through daily antivirus content updates (active Threat Prevention license required).

STEP 2 | Find an external dynamic list to use with the firewall.

- Create an external dynamic list and host it on a web server. Enter IP addresses, domains, or URLs in a blank text file. Each list entry must be on a separate line. For example:

  ```
  financialtimes.co.in

  www.wallaby.au/joey

  www.exyang.com/auto-tutorials/How-to-enter-Data-for-Success.aspx

  *.example.com/*

  abc?*/abc.com

  *&*.net
  ```

  See the Formatting Guidelines for an External Dynamic List to ensure that the firewall does not skip list entries. To prevent commit errors and invalid entries, do not prefix http:// or https:// to any of the entries.
- Use an external dynamic list hosted by another source and verify that it follows the Formatting Guidelines for an External Dynamic List.

STEP 3 | Select **Objects** > **External Dynamic Lists**.

STEP 4 | Click **Add** and enter a descriptive **Name** for the list.

STEP 5 | (Optional) Select **Shared** to share the list with all virtual systems on a device that is enabled for multiple virtual systems. By default, the object is created on the virtual system that is currently selected in the **Virtual Systems** drop-down.

STEP 6 | (Panorama only) Select **Disable override** to ensure that a firewall administrator cannot override settings locally on a firewall that inherits this configuration through a Device Group commit from Panorama.

STEP 7 | Select the list **Type** (for example, **URL List**).

Ensure that the list only includes entries for the list type. See Verify whether entries in the external dynamic list were ignored or skipped.

STEP 8 | Enter the **Source** for the list you just created on the web server. The source must include the full path to access the list. For example, `https://1.2.3.4/EDL_IP_2015`.

If you are creating a list of type Predefined IP, select a Palo Alto Networks malicious IP address feed to use as a source.

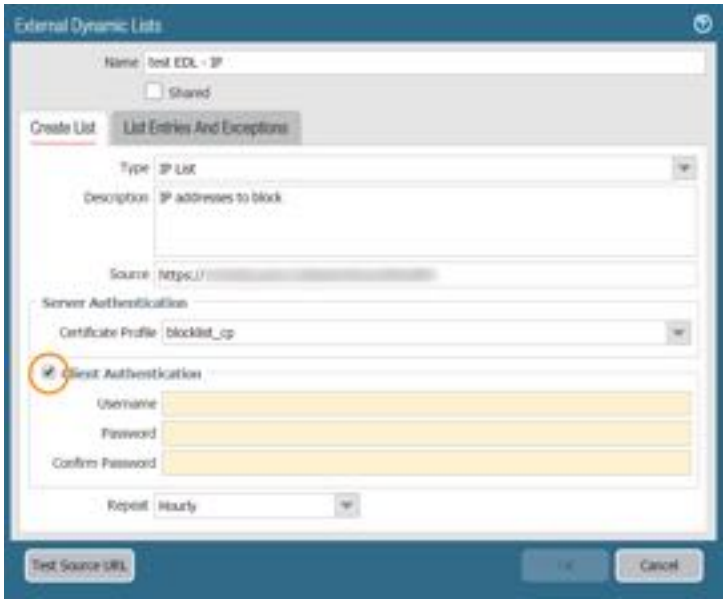
© 2018 Palo Alto Networks, Inc.

STEP 9 | If the list source is secured with SSL (i.e. lists with an HTTPS URL), enable server authentication. Select a **Certificate Profile** or create a **New Certificate Profile** for authenticating the server that hosts the list. The certificate profile you select must have root CA (certificate authority) and intermediate CA certificates that match the certificates installed on the server you are authenticating.

Maximize the number of external dynamic lists that you can use to enforce policy. Use the same certificate profile to authenticate external dynamic lists from the same source URL. If you assign different certificate profiles to external dynamic lists from the same source URL, the firewall counts each list as a unique external dynamic list.

STEP 10 | Enable client authentication if the list source has an HTTPS URL and requires basic HTTP authentication for list access.

1. Select **Client Authentication**.
2. Enter a valid **Username** to access the list.
3. Enter the **Password** and **Confirm Password**.



STEP 11 | (Not available on Panorama) Click **Test Source URL** to verify that the firewall can connect to the web server.

STEP 12 | (Optional) Specify the **Repeat** frequency at which the firewall retrieves the list. By default, the firewall retrieves the list once every hour and commits the changes.

The interval is relative to the last commit. So, for the five-minute interval, the commit occurs in 5 minutes if the last commit was an hour ago. To retrieve the list immediately, see Retrieve an External Dynamic List from the Web Server.

STEP 13 | Click **OK** and **Commit**.

STEP 14 | Enforce Policy on an External Dynamic List.

If the server or client authentication fails, the firewall ceases to enforce policy based on the last successfully retrieved external dynamic list. Find External Dynamic Lists That Failed Authentication *and view the reasons for authentication failure.*

© 2018 Palo Alto Networks, Inc.

Retrieve an External Dynamic List from the Web Server

When you Configure the Firewall to Access an External Dynamic List, you can configure the firewall to retrieve the list from the web server on an hourly (default)five minute, daily, weekly, or monthly basis. If you have added or deleted IP addresses from the list and need to trigger an immediate refresh, use the following process to fetch the updated list.

STEP 1 | To retrieve the list on demand, select **Objects** > **External Dynamic Lists**.

STEP 2 | Select the list that you want to refresh, and click **Import Now**. The job to import the list is queued.

STEP 3 | To view the status of the job in the Task Manager, see Manage and Monitor Administrative Tasks.

STEP 4 | (Optional) After the firewall retrieves the list, View External Dynamic List Entries.

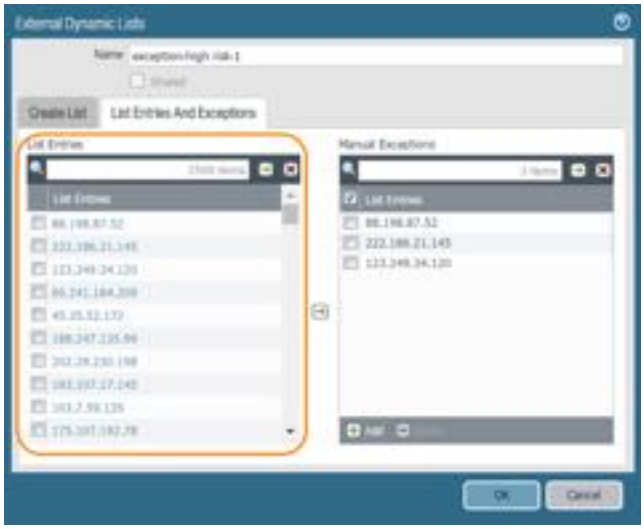
View External Dynamic List Entries

Before you Enforce Policy on an External Dynamic List, you can view the contents of an external dynamic list directly on the firewall to check if it contains certain IP addresses, domains, or URLs. The entries displayed are based on the version of the external dynamic list that the firewall most recently retrieved.

STEP 1 | Select **Objects** > **External Dynamic Lists**.

STEP 2 | Click the external dynamic list you want to view.

STEP 3 | Click **List Entries and Exceptions** and view the objects that the firewall retrieved from the list.



The list might be empty if:

- The firewall has not yet retrieved the external dynamic list. To force the firewall to retrieve an external dynamic list immediately, Retrieve an External Dynamic List from the Web Server.
- The firewall is unable to access the server that hosts the external dynamic list. Click **Test Source URL** to verify that the firewall can connect to the server.

© 2018 Palo Alto Networks, Inc.

STEP 4 | Enter an IP address, domain, or URL (depending on the type of list) in the filter field and Apply Filter ( ) to check if it's in the list. Exclude Entries from an External Dynamic List based on which IP addresses, domains, and URLs you need to block or allow.

STEP 5 | (Optional) View the AutoFocus Intelligence Summary for a list entry. Hover over an entry to open the drop-down and then click **AutoFocus**.

Exclude Entries from an External Dynamic List

As you view the entries of an external dynamic list, you can exclude up to 100 entries from the list. The ability to exclude entries from an external dynamic list gives you the option to enforce policy on some (but not all) of the entries in a list. This is helpful if you cannot edit the contents of an external dynamic list (such as the Palo Alto Networks High-Risk IP Addresses feed) because it comes from a third-party source.

STEP 1 | View External Dynamic List Entries.

STEP 2 | Select up to 100 entries to exclude from the list and click Submit ( ) or manually **Add** a list exception.

- You cannot save your changes to the external dynamic list if you have duplicate entries in the Manual Exceptions list. To identify duplicate entries, look for entries with a red underline.
- A manual exception must match a list entry exactly. For example, if an IP address range is included as a list entry and you manually enter a single IP address within the range as a list exception, the firewall will continue to enforce policy on all the IP addresses in the range. So, to exclude that single IP address, you must first make sure that it's a standalone external dynamic list entry, and then manually add the same IP address to the list of exceptions.

STEP 3 | Click **OK** and **Commit** to save your changes.

STEP 4 | (Optional) Enforce Policy on an External Dynamic List.

Enforce Policy on an External Dynamic List

Block or allow traffic based on IP addresses or URLs in an external dynamic list, or use an dynamic domain list with a DNS sinkhole to prevent access to malicious domains. Refer to the table below for the ways you can use external dynamic lists to enforce policy on the firewall.

- Configure DNS Sinkholing for a List of Custom Domains.
- Use an External Dynamic List in a URL Filtering Profile.
- Use an External Dynamic List of Type URL as Match Criteria in a Security Policy Rule.
  1. Select **Policies** > **Security**.
  2. Click **Add** and enter a descriptive **Name** for the rule.
  3. In the **Source** tab, select the **Source Zone**.
  4. In the **Destination** tab, select the **Destination Zone**.
  5. In the **Service/URL Category** tab, click **Add** to select the appropriate external dynamic list from the URL Category list.
  6. In the **Actions** tab, set the **Action Setting** to **Allow** or **Deny**.
  7. Click **OK** and **Commit**.
  8. Verify whether entries in the external dynamic list were ignored or skipped.

     Use the following CLI command on a firewall to review the details for a list.

© 2018 Palo Alto Networks, Inc.

```
request system external-list show type <domain | ip | url> name_of_list
```

For example:

```
request system external-list show type url EBL_ISAC_Alert_List
```

9. Test that the policy action is enforced.
   1. View External Dynamic List Entries for the URL list, and attempt to access a URL from the list.
   2. Verify that the action you defined is enforced.
   3. To monitor the activity on the firewall:
      - Select **ACC** and add a URL Domain as a global filter to view the Network Activity and Blocked Activity for the URL you accessed.
      - Select **Monitor** > **Logs** > **URL Filtering** to access the detailed log view.

- Use an External Dynamic List of Type IP or Predefined IP as a Source or Destination Address Object in a Security Policy Rule.

  This capability is useful if you deploy new servers and want to allow access to the newly deployed servers without requiring a firewall commit.

  1. Select **Policies** > **Security**.
  2. Click Add and give the rule a descriptive name in the General tab.
  3. In the Source tab, select the Source Zone and optionally select the external dynamic list as the Source Address.
  4. In the Destination tab, select the Destination Zone and optionally select the external dynamic list as the Destination Address.
  5. In the Service/ URL Category tab, make sure the Service is set to application-default.
  6. In the Actions tab, set the Action Setting to **Allow** or **Deny**.

     Create separate external dynamic lists if you want to specify allow and deny actions for specific IP addresses.
  7. Leave all the other options at the default values.
  8. Click **OK** to save the changes.
  9. **Commit** the changes.
  10. Test that the policy action is enforced.
      1. View External Dynamic List Entries for the external dynamic list, and attempt to access an IP address from the list.
      2. Verify that the action you defined is enforced.
      3. Select **Monitor** > **Logs** > **Traffic** and view the log entry for the session.
      4. To verify the policy rule that matches a flow, use the following CLI command:

```
test security-policy-match source <IP_address> destination <IP_address>
 destination port <port_number> protocol <protocol_number>
```

      Tips for enforcing policy on the firewall with external dynamic lists:
      - *When viewing external dynamic lists on the firewall (Objects > External Dynamic Lists), click List Capacities to compare how many IP addresses, domains, and URLs are currently used in policy with the total number of entries that the firewall supports for each list type.*

© 2018 Palo Alto Networks, Inc.

- Use Global Find to Search the Firewall or Panorama Management Server *for a domain, IP address, or URL that belongs to one or more external dynamic lists is used in policy. This is useful for determining which external dynamic list (referenced in a Security policy rule) is causing the firewall to block or allow a certain domain, IP address, or URL.*

Find External Dynamic Lists That Failed Authentication

When an external dynamic list that requires SSL fails client or server authentication, the firewall generates a system log of critical severity. The log is critical because the firewall ceases to enforce policy based on the external dynamic list after it fails authentication. Use the following process to view critical system log messages notifying you of authentication failure related to external dynamic lists.

STEP 1 | Select **Monitor** > **Logs** > **System**.

STEP 2 | Construct the following filters to view all messages related to authentication failure, and apply the filters. For more information, review the complete workflow to Filter Logs.

- Server authentication failure—**`(eventid eq tls-edl-auth-failure)`**
- Client authentication failure—**`(eventid eq edl-cli-auth-failure)`**



STEP 3 | Review the system log messages. The message description includes the name of the external dynamic list, the source URL for the list, and the reason for the authentication failure.

The server that hosts the external dynamic list fails authentication if the certificate is expired. If you have configured the certificate profile to check certificate revocation status via Certificate Revocation List (CRL) or Online Certificate Status Protocol (OCSP), the server may also fail authentication if:

- The certificate is revoked.
- The revocation status of the certificate is unknown.
- The connection times out as the firewall is attempting to connect to the CRL/OCSP service.

For more information on certificate profile settings, refer to the steps to Configure a Certificate Profile.

Verify that you added the root CA and intermediate CA of the server to the certificate profile configured with the external dynamic list. Otherwise, the firewall will not authenticate the list properly.

Client authentication fails if you have entered the incorrect username and password combination for the external dynamic list.

STEP 4 | (Optional) Disable Authentication for an External Dynamic List that failed authentication as a stop-gap measure until the list owner renews the certificate(s) of the server that hosts the list.

© 2018 Palo Alto Networks, Inc.

Disable Authentication for an External Dynamic List

Palo Alto Networks recommends that you enable authentication for the servers that host the external dynamic lists configured on your firewall. However, if you Find External Dynamic Lists That Failed Authentication and prefer to disable server authentication for those lists, you can do so through the CLI. The procedure below only applies to external dynamic lists secured with SSL (i.e., lists with an HTTPS URL); the firewall does not enforce server authentication on lists with an HTTP URL.



Disabling server authentication for an external dynamic list also disables client authentication. With client authentication disabled, the firewall will not be able to connect to an external dynamic list that requires a username and password for access.

STEP 1 | Launch the CLI and switch to configuration mode as follows:

```
username@hostname> configure
Entering configuration mode
[edit]
username@hostname#
```

The change from the > to the # symbol indicates that you are now in configuration mode.

STEP 2 | Enter the appropriate CLI command for the list type:

- IP Address

```
set external-list <external dynamic list name> type ip certificate-profile
 None
```

- Domain

```
set external-list <external dynamic list name> type domain certificate-
profile None
```

- URL

```
set external-list <external dynamic list name> type url certificate-
profile None
```

STEP 3 | Verify that authentication is disabled for the external dynamic list.

Trigger a refresh for the list (see Retrieve an External Dynamic List from the Web Server). If the firewall retrieves the list successfully, server authentication is disabled.

© 2018 Palo Alto Networks, Inc.

Register IP Addresses and Tags Dynamically

To mitigate the challenges of scale, lack of flexibility and performance, the architecture in networks today allows for clients, servers, and applications to be provisioned, changed, and deleted on demand. This agility poses a challenge for security administrators because they have limited visibility into the IP addresses of the dynamically provisioned clients and servers, and the plethora of applications that can be enabled on these virtual resources.

The firewall (hardware-based models and the VM-Series) supports the ability to register IP addresses and tags dynamically. The IP addresses and tags can be registered on the firewall directly or registered on the firewall through Panorama. You can also automatically remove tags on the source or destination IP address included in a firewall log.

This dynamic registration process can be enabled using any of the following options:

- **User-ID agent for Windows**—In an environment where you've deployed the User-ID agent, you can enable the User-ID agent to monitor up to 100 VMware ESXi and/or vCenter Servers. As you provision or modify virtual machines on these VMware servers, the agent can retrieve the IP address changes and share them with the firewall.
- **VM Information Sources**—Allows you to monitor VMware ESXi and vCenter Server and the AWS-VPC natively on the firewall, to retrieve IP address changes when you provision or modify virtual machines on these sources. For monitoring virtual machines in your Microsoft Azure deployment, you must deploy the VM Monitoring script that runs on a virtual machine within the Azure public cloud. This script collects the IP address-to-tag mapping for all your Azure assets and uses the API to push the VM information to your Palo Alto Networks® firewall(s). VM Information Sources polls for a predefined set of attributes and does not require external scripts to register the IP addresses through the XML API. See Monitor Changes in the Virtual Environment.
- **VMware Service Manager** (only available for the integrated NSX solution)—The integrated NSX solution is designed for automated provisioning and distribution of Palo Alto Networks next-generation security services and the delivery of dynamic context-based security policies using Panorama. The NSX Manager updates Panorama with the latest information on the IP addresses and tags associated with the virtual machines deployed in this integrated solution. For information on this solution, see Set Up a VM-Series NSX Edition Firewall.
- **XML API**—The firewall and Panorama support an XML API that uses standard HTTP requests to send and receive data. You can use this API to register IP addresses and tags with the firewall or Panorama. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports REST-based services. Refer to the PAN-OS XML API Usage Guide for details.
- **Auto-Tag**— Tag the source or destination IP address automatically when a log is generated on the firewall, and register the IP address and tag mapping to a User-ID agent on the firewall or Panorama, or to a remote User-ID agent using an HTTP server profile. For example, whenever the firewall generates a threat log, you can configure the firewall to tag the source IP address in the threat log with a specific tag name. See Forward Logs to an HTTP(S) Destination.

For information on creating and using Dynamic Address Groups, see Use Dynamic Address Groups in Policy.

For the CLI commands for registering tags dynamically, see CLI Commands for Dynamic IP Addresses and Tags.

© 2018 Palo Alto Networks, Inc.

Monitor Changes in the Virtual Environment

To secure applications and prevent threats in an environment where new users and servers are constantly emerging, your security policy must be nimble. To be nimble, the firewall must be able to learn about new or modified IP addresses and consistently apply policy without requiring configuration changes on the firewall.

This capability is provided by the coordination between the **VM Information Sources** and **Dynamic Address Groups** features on the firewall. The firewall and Panorama provide an automated way to gather information on the virtual machine (or guest) inventory on each monitored source and create policy objects that stay in sync with the dynamic changes on the network.

- Enable VM Monitoring to Track Changes on the Virtual Network
- Attributes Monitored on Virtual Machines in Cloud Platforms
- Use Dynamic Address Groups in Policy

Enable VM Monitoring to Track Changes on the Virtual Network

VM information sources provides an automated way to gather information on the Virtual Machine (VM) inventory on each monitored source (host); the firewall can monitor the VMware ESXi and vCenter Server, the AWS-VPC and Azure VNet. As virtual machines (guests) are deployed or moved, the firewall collects a predefined set of attributes (or metadata elements) as tags; these tags can then be used to define Dynamic Address Groups (see Use Dynamic Address Groups in Policy) and matched against in policy.

For monitoring the virtual machine on VMware ESXi, vCenter Server, and the AWS-VPC, you can directly configure the firewall or use Panorama templates to monitor up to 10 VM information sources. **VM Information Sources** offers easy configuration and enables you to monitor a predefined set of 16 metadata elements or attributes. See Attributes Monitored on Virtual Machines in Cloud Platforms for the list. By default, the traffic between the firewall and the monitored sources uses the management (MGT) port on the firewall.



- *When monitoring ESXi hosts that are part of the VM-Series NSX edition solution, use Dynamic Address Groups instead of using VM Information Sources to learn about changes in the virtual environment. For the VM-Series NSX edition solution, the NSX Manager provides Panorama with information on the NSX security group to which an IP address belongs. The information from the NSX Manager provides the full context for defining the match criteria in a Dynamic Address Group because it uses the service profile ID as a distinguishing attribute and allows you to properly enforce policy when you have overlapping IP addresses across different NSX security groups. Up to a maximum of 32 tags (from vCenter server and NSX Manager) that can be registered to an IP address.*
- *For monitoring the virtual machines within your Azure deployment, instead of VM Monitoring Sources, you need to deploy the VM Monitoring script that runs on a virtual machine within the Azure public cloud. This script collects the IP address-to-tag mapping information for your Azure assets and publishes it to the firewalls and corresponding virtual systems you specify in the script.*

STEP 1 | Enable VM Monitoring.



You can configure up to 10 VM information sources for each firewall, or for each virtual system on a multiple virtual systems capable firewall.

If your firewalls are configured in a high availability configuration:

© 2018 Palo Alto Networks, Inc.

- In an active/passive setup, only the active firewall monitors the VM sources.
- In an active/active setup, only the firewall with the priority value of primary monitors the VM sources.

1. Select **Device** > **VM Information Sources**.
2. Click **Add** and enter the following information:
   - A **Name** to identify the source that you want to monitor.
   - Enter the **Host information for the server—**hostname or IP address and the **Port** on which it is listening.
   - Select the **Type** to indicate whether the source is an **AWS VPC**, a **VMware ESX(i)** server or a **VMware vCenter** Server.
   - Add the credentials (**Username** and **Password**) to authenticate to the server specified above.
   - Use the credentials of an administrative user to enable access.
   - (Optional) Modify the **Update interval** to a value between 5-600 seconds. By default, the firewall polls every 5 seconds. The API calls are queued and retrieved within every 60 seconds, so updates may take up to 60 seconds plus the configured polling interval.



   - (Optional) Enter the interval in hours when the connection to the monitored source is closed, if the host does not respond. (range is 2-10 hours; default is 2).

     To change the default value, select the check box to **Enable timeout when the source is disconnected** and specify the value. When the specified limit is reached or if the host cannot be accessed or does not respond, the firewall will close the connection to the source.
   - Click **OK**, and **Commit** the changes.
   - Verify that the connection **Status** displays as connected.

STEP 2 | Verify the connection status.

Verify that the connection **Status** displays as connected.



If the connection status is pending or disconnected, verify that the source is operational and that the firewall is able to access the source. If you use a port other than the MGT port for communicating with

© 2018 Palo Alto Networks, Inc.

the monitored source, you must change the service route (**Device** > **Setup** > **Services**, click the **Service Route Configuration** link and modify the **Source Interface** for the **VM Monitor** service).

Attributes Monitored on Virtual Machines in Cloud Platforms

In order to collect the values assigned to the monitored VMs, the next-gen firewall monitors the following predefined set of attributes on the VMware ESXi/vCenter Server, Amazon Web Services (AWS)-VPC, and within your Microsoft Azure VNet deployment:

Attributes Monitored on a VMware Source	Attributes Monitored on the AWS-VPC	Attributes Monitored on Microsoft Azure
UUID	Architecture	VM Name
Name	Guest OS	VM Size
Guest OS	Image ID	OS Type
VM State — the power state can be poweredOff, poweredOn, standBy, and unknown.	Instance ID	OS Publisher
Annotation	Instance State	OS Offer
Version	Instance Type	OS SKU
Network — Virtual Switch Name, Port Group Name, and VLAN ID	Key Name	Subnet
Container Name —vCenter Name, Data Center Object Name, Resource Pool Name, Cluster Name, Host, Host IP address.	• Placement—Tenancy, Group Name, Availability Zone • Private DNS Name • Public DNS Name • Subnet ID • Tag (key, value); up to 18 tags supported per instance • VPC ID	VNet
		Azure Region
		Resource Group Name
		User Defined Tags

Use Dynamic Address Groups in Policy

Dynamic address groups are used in policy. They allow you to create policy that automatically adapts to changes—adds, moves, or deletions of servers. It also enables the flexibility to apply different rules to the

© 2018 Palo Alto Networks, Inc.

same server based on *tags* that define its role on the network, the operating system, or the different kinds of traffic it processes.

A dynamic address group uses tags as a filtering criteria to determine its members. The filter uses logical *and* and *or* operators. All IP addresses or address groups that match the filtering criteria become members of the dynamic address group. Tags can be defined statically on the firewall and/or registered (dynamically) to the firewall. The difference between static and dynamic tags is that static tags are part of the configuration on the firewall, and dynamic tags are part of the runtime configuration. This implies that a commit is not required to update dynamic tags; the tags must however be used by Dynamic Address Groups that are referenced in policy, and the policy must be committed on the firewall.

To dynamically register tags, you can use the XML API or the VM Monitoring agent on the firewall or on the User-ID agent. Each tag is a metadata element or attribute-value pair that is registered on the firewall or Panorama. For example, IP1 {tag1, tag2,.....tag32}, where the IP address and the associated tags are maintained as a list; each registered IP address can have up to 32 tags such as the operating system, the datacenter or the virtual switch to which it belongs. Within 60 seconds of the API call, the firewall registers the IP address and associated tags, and automatically updates the membership information for the dynamic address group(s).

The maximum number of IP addresses that can be registered for each model is different. Use the following table for specifics on your model:

Model	Maximum number of dynamically registered IP addresses
PA-7000 Series, PA-5060, PA-5200 Series, VM-300, VM-500, VM-700, VM-1000-HV	100,000
PA-5050	50,000
PA-5020	25,000
PA-3000 Series	5,000
PA-850, VM-100	2,500
PA-820, PA-500, PA-220, PA-200 VM-200, VM-50	1000

The following example shows how dynamic address groups can simplify network security enforcement. The example workflow shows how to:

- Enable the VM Monitoring agent on the firewall, to monitor the VMware ESX(i) host or vCenter Server and register VM IP addresses and the associated tags.
- Create dynamic address groups and define the tags to filter. In this example, two address groups are created. One that only filters for dynamic tags and another that filters for both static and dynamic tags to populate the members of the group.
- Validate that the members of the dynamic address group are populated on the firewall.
- Use dynamic address groups in policy. This example uses two different security policies:
  - A security policy for all Linux servers that are deployed as FTP servers; this rule matches on dynamically registered tags.

© 2018 Palo Alto Networks, Inc.

- A security policy for all Linux servers that are deployed as web servers; this rule matches on a dynamic address group that uses static and dynamic tags.

- Validate that the members of the dynamic address groups are updated as new FTP or web servers are deployed. This ensure that the security rules are enforced on these new virtual machines too.

STEP 1 | Enable VM Source Monitoring.

See Enable VM Monitoring to Track Changes on the Virtual Network.

STEP 2 | Create dynamic address groups on the firewall.

View the tutorial to see a big picture view of the feature.

1. Log in to the web interface of the firewall.
2. Select **Object** > **Address Groups**.
3. Click **Add** and enter a **Name** and a **Description** for the address group.
4. Select **Type** as **Dynamic**.
5. Define the match criteria. You can select dynamic and static tags as the match criteria to populate the members of the group. Click **Add Match Criteria**, and select the **And** or **Or** operator and select the attributes that you would like to filter for or match against. and then click **OK**.



6. Click **Commit**.

STEP 3 | The match criteria for each dynamic address group in this example is as follows:

ftp_server: matches on the guest operating system “Linux 64-bit” and annotated as “ftp” ('guestos.Ubuntu Linux 64-bit' and 'annotation.ftp').

web-servers: matches on two criteria—the tag black or if the guest operating system is Linux 64-bit and the name of the server us Web_server_Corp. ('guestos.Ubuntu Linux 64-bit' and 'vmname.WebServer_Corp' or 'black')



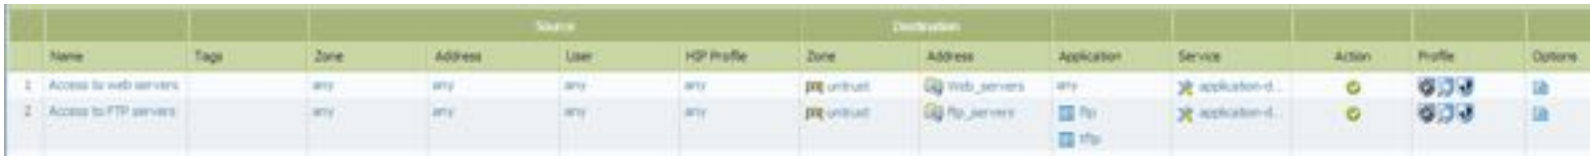
STEP 4 | Use dynamic address groups in policy.

View the tutorial.

© 2018 Palo Alto Networks, Inc.

1. Select **Policies** > **Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the Dynamic address group you just created.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeats Steps 1 through 6 to create another policy rule.
8. Click **Commit**.

STEP 5 | This example shows how to create two policies: one for all access to FTP servers and the other for access to web servers.

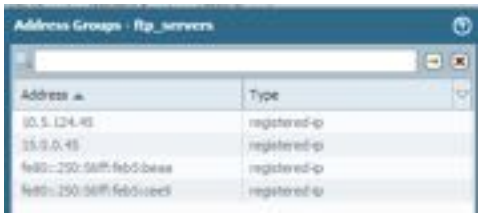


STEP 6 | Validate that the members of the dynamic address group are populated on the firewall.

1. Select **Policies** > **Security**, and select the rule.
2. Select the drop-down arrow next to the address group link, and select **Inspect**. You can also verify that the match criteria is accurate.



3. Click the **more** link and verify that the list of registered IP addresses is displayed.



Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.

© 2018 Palo Alto Networks, Inc.

CLI Commands for Dynamic IP Addresses and Tags

The Command Line Interface on the firewall and Panorama give you a detailed view into the different sources from which tags and IP addresses are dynamically registered. It also allows you to audit registered and unregistered tags. The following examples illustrate the capabilities in the CLI.

Example	CLI Command
View all registered IP addresses that match the tag, `state.poweredOn` or that are not tagged as `vSwitch0`.	**show log iptag tag_name equal state.poweredOn** **show log iptag tag_name not-equal switch.vSwitch0**
View all dynamically registered IP addresses that were sourced by VM Information Source with name `vmware1` and tagged as `poweredOn`.	**show vm-monitor source source-name vmware1 tag state.poweredOn registered-ip all** registered IP Tags ---------- ---------- fe80::20c:29ff:fe69:2f76 "state.poweredOn" 10.1.22.100 "state.poweredOn" 2001:1890:12f2:11:20c:29ff:fe69:2f76"state.poweredOn" fe80::20c:29ff:fe69:2f80 "state.poweredOn" 192.168.1.102 "state.poweredOn" 10.1.22.105 "state.poweredOn" 2001:1890:12f2:11:2cf8:77a9:5435:c0d"state.poweredOn" fe80::2cf8:77a9:5435:c0d "state.poweredOn"
Clear all IP addresses and tags learned from a specific VM Monitoring source without disconnecting the source.	**debug vm-monitor clear source-name *<name>***
Display IP addresses registered from all sources.	**show object registered-ip all**
Display the count for IP addresses registered from all sources.	**show object registered-ip all option count**
Clear IP addresses registered from all sources	**debug object registered-ip clear all**
Add or delete tags for a given IP address that was registered using the XML API.	**debug object registered-ip test [<register/unregister>] <ip/netmask><tag>**
View all tags registered from a specific information source.	**show vm-monitor source source-name vmware1 tag all**

© 2018 Palo Alto Networks, Inc.

Example	CLI Command
	vlanId.4095 vswitch.vSwitch1 host-ip.10.1.5.22 portgroup.TOBEUSED hostname.panserver22 portgroup.VM Network 2 datacenter.ha-datacenter vlanId.0 state.poweredOn vswitch.vSwitch0 vmname.Ubuntu22-100 vmname.win2k8-22-105 resource-pool.Resources vswitch.vSwitch2 guestos.Ubuntu Linux 32-bit guestos.Microsoft Windows Server 2008 32-bit annotation. version.vmx-08 portgroup.VM Network vm-info-source.vmware1 uuid.564d362c-11cd-b27f-271f-c361604dfad7 uuid.564dd337-677a-eb8d-47db-293bd6692f76 Total: 22
View all tags registered from a specific data source, for example from the VM Monitoring Agent on the firewall, the XML API, Windows User-ID Agent or the CLI.	• To view tags registered from the CLI: **show log iptag datasource_type equal unknown** • To view tags registered from the XML API: **show log iptag datasource_type equal xml-api** • To view tags registered from VM Information sources: **show log iptag datasource_type equal vm-monitor** • To view tags registered from the Windows User-ID agent: **show log iptag datasource_type equal xml-api datasource_subtype equal user-id-agent**
View all tags that are registered for a specific IP address (across all sources).	**debug object registered-ip show tag-source ip ip_address tag all**

© 2018 Palo Alto Networks, Inc.

Identify Users Connected through a Proxy Server

If you have a proxy server deployed between the users on your network and the firewall, the firewall might see the proxy server IP address as the source IP address in HTTP/HTTPS traffic that the proxy forwards rather than the IP address of the client that requested the content. In many cases, the proxy server adds an X-Forwarded-For (XFF) header to HTTP requests that include the actual IPv4 or IPv6 address of the client that requested the content or from whom the request originated. In such cases, you can configure the firewall to extract the end user IP address from the XFF so that User-ID can map that IP address to the username of the end user. This enables you enforce user-based policy to safely enable access to web-based applications for your users behind a proxy server. In addition, if User-ID is able to map the XFF IP address to a username, the firewall displays that username as the Source user in Traffic, Threat, WildFire Submissions, and URL Filtering logs for visibility into the web activity of users behind the proxy.

To use the XFF header for user mapping:

- The XFF header your proxy server adds must contain the source IP address of the end user who originated the request. If the header contains multiple IP addresses, the firewall uses the first IP address only. If the header contains information other than an IP address, the firewall will not be able to perform user mapping.
- You must Enable User-ID.

With this option enabled, the firewall uses the IP address in the XFF header for user mapping purposes only. The source IP address the firewall logs is still that of the proxy server, not that of the source user. When you see a log event attributed to a user that the firewall mapped using and IP address extracted from an XFF header, it can be difficult to track down the specific device associated with the event. To simplify debugging and troubleshooting of events attributed to users behind the proxy server, you must also configure the firewall to populate the X-Forwarded-For column in the URL Filtering log with the IP address in the XFF header so that you can track down the specific user and device associated with an log event that is correlated with the URL Filtering log entry.

- Use XFF Values for Policies and Logging Source Users
- Use the IP Address in the XFF Header to Troubleshoot Events

Use XFF Values for Policies and Logging Source Users

You can configure the firewall map the IP address in the XFF header to a username using User-ID so that you can have visibility into and user-based policy control over the web traffic of users behind a proxy server who cannot otherwise be identified. In order to map the IP addresses from the XFF headers to usernames, you must first Enable User-ID.



Enabling the firewall to use the X-Forwarded-For headers to perform user mapping does not enable the firewall to use the client IP address in the XFF header as the source address in the logs; the logs still display the proxy server IP address as the source address. However, to simplify the debugging and troubleshooting process you can configure the firewall to Use the IP Address in the XFF Header to Troubleshoot Events *to display the client IP address from the XFF header in the URL Filtering logs.*

To ensure that attackers can't read and exploit the XFF values in web request packets that exit the firewall to retrieve content from an external server, you can also configure the firewall to strip the XFF values from outgoing packets.

© 2018 Palo Alto Networks, Inc.

These options are not mutually exclusive: if you configure both, the firewall zeroes out XFF values only after using them in policy enforcement and logging.

STEP 1 | Enable the firewall to use XFF values in policies and in the source user fields of logs.

1. Select **Device** > **Setup** > **Content-ID** and edit the X-Forwarded-For Headers settings.
2. Select **Use X-Forwarded-For Header in User-ID**.

STEP 2 | Remove XFF values from outgoing web requests.

1. Select **Strip X-Forwarded-For Header**.
2. Click **OK** and **Commit**.

STEP 3 | Verify the firewall is populating the source user fields of logs.

1. Select a log type that has a source user field (for example, **Monitor** > **Logs** > **Traffic**).
2. Verify that the Source User column displays the usernames of users who access the web.

Use the IP Address in the XFF Header to Troubleshoot Events

By default, the firewall does not log the source address of a client behind a proxy server, even if you are using this address from the X-Forwarded-For (XFF) header for user mapping. Therefore, while you can identify the specific user associated with a log event, you will not be able to easily identify the source device that originated the log event. To simplify the debugging and troubleshooting of events for users behind a proxy server, you must enable the X-Forwarded-For option within HTTP Header Logging in the URL Filtering profile that you attach to security policy rules that allow access to web-based applications. With this option enabled, the firewall logs the IP address from the XFF header as the Source address for all traffic that matches the rule.

Enabling the firewall to use the XFF header as the Source address in URL Filtering logs does not enable user mapping of the source address. To populate the source user fields, see Use XFF Values for Policies and Logging Source Users.

STEP 1 | Enable the X-Forwarded-For option within HTTP Header Logging in the URL Filtering profile.

1. Select **Objects** > **Security Profiles** > **URL Filtering** and select the URL Filtering profile you want to configure, or add a new one.

   You can't enable XFF logging in the default URL Filtering profile.

2. Select the **Settings** tab and select **X-Forwarded-For**.
3. Click **OK** to save the profile.

STEP 2 | Attach the URL Filtering profile to the security policy rule(s) that enable access to web applications.

1. Select **Policies** > **Security** and click the rule.
2. Select the **Actions** tab, set the **Profile Type** to **Profiles**, and select the **URL Filtering** profile you just configured for X-Forwarded-For HTTP Header Logging.
3. Click **OK** and **Commit**.

STEP 3 | Verify the firewall is logging XFF values.

1. Select **Monitor** > **Logs** > **URL Filtering**.
2. View the XFF values in one of the following ways:

© 2018 Palo Alto Networks, Inc.

- To display the XFF value for a single URL Filtering log—Click the spyglass icon for the log to displays its details. The HTTP Headers section displays the X-Forwarded-For value.
- To display the XFF values for all URL Filtering logs—Open the drop-down in any column header, select **Columns**, and select **X-Forwarded-For**. The page then displays an X-Forwarded-For column.

STEP 4 | Use the XFF field in the URL Filtering log to troubleshoot a log event in another log type.

Although only the URL Filtering logs display the IP address of the source user in the X-Forwarded-For column of the logs, if you notice an event associated with HTTP/HTTPS traffic but that you cannot identify the source IP address because it is that of the proxy server, you can use the X-Forwarded-For value in a correlated URL Filtering log to help you identify the source address associated with the log event. To do this:

1. Find an event you want investigate in a Traffic, Threat, or WildFire Submissions logs that is showing the IP address of the proxy server as the source address.
2. Click the spyglass icon for the log to display its details and look for an associated URL Filtering log at the bottom of the Detailed Log Viewer window.
3. Select the header row and then select **X-Forwarded-For** from the **Columns** drop-down to display this value. The IP address in this column of the X-Forwarded-For column represents the IP address of the source user behind the proxy server. Use this IP address to track down the device that triggered the event you are investigating.

© 2018 Palo Alto Networks, Inc.

Policy-Based Forwarding

Normally, the firewall uses the destination IP address in a packet to determine the outgoing 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, and specify the outgoing or *egress* interface based on specific parameters such as source or destination IP address, or type of traffic.

- PBF on page 1033
- Create a Policy-Based Forwarding Rule on page 1035
- Use Case: PBF for Outbound Access with Dual ISPs on page 1036

PBF

PBF rules allow traffic to take an alternative path from the next hop specified in the route table, and are typically used to specify an egress interface for security or performance reasons. Let's say your company has two links between the corporate office and the branch office: a cheaper internet link and a more expensive leased line. The leased line is a high-bandwidth, low-latency link. For enhanced security, you can use PBF to send applications that aren't encrypted traffic, such as FTP traffic, over the private leased line and all other traffic over the internet link. Or, for performance, you can choose to route business-critical applications over the leased line while sending all other traffic, such as web browsing, over the cheaper link.

- Egress Path and Symmetric Return on page 1033
- Path Monitoring for PBF on page 1033
- Service Versus Applications in PBF on page 1034

Egress Path and Symmetric Return

Using PBF, you can direct traffic to a specific interface on the firewall, drop the traffic, or direct traffic to another virtual system (on systems enabled for multiple virtual systems).

In networks with asymmetric routes, such as in a dual ISP environment, connectivity issues occur when traffic arrives at one interface on the firewall and leaves from another interface. If the route is asymmetrical, where the forward (SYN packet) and return (SYN/ACK) paths are different, the firewall is unable to track the state of the entire session and this causes a connection failure. To ensure that the traffic uses a symmetrical path, which means that the traffic arrives at and leaves from the same interface on which the session was created, you can enable the *Symmetric Return* option.

With symmetric return, the virtual router overrides a routing lookup for return traffic and instead directs the flow back to the MAC address from which it received the SYN packet (or first packet). However, if the destination IP address is on the same subnet as the ingress/egress interface's IP address, a route lookup is performed and symmetric return is not enforced. This behavior prevents traffic from being blackholed.



To determine the next hop for symmetric returns, the firewall uses an Address Resolution Protocol (ARP) table. The maximum number of entries that this ARP table supports is limited by the firewall model and the value is not user configurable. To determine the limit for your model, use the CLI command: ***`show pbf return-mac all`***.

Path Monitoring for PBF

Path monitoring allows you to verify connectivity to an IP address so that the firewall can direct traffic through an alternate route, when needed. The firewall uses ICMP pings as *heartbeats* to verify that the specified IP address is reachable.

© 2018 Palo Alto Networks, Inc.

A monitoring profile allows you to specify the threshold number of heartbeats to determine whether the IP address is reachable. When the monitored IP address is unreachable, you can either disable the PBF rule or specify a *fail-over* or *wait-recover* action. Disabling the PBF rule allows the virtual router to take over the routing decisions. When the fail-over or wait-recover action is taken, the monitoring profile continues to monitor whether the target IP address is reachable, and when it comes back up, the firewall reverts back to using the original route.

The following table lists the difference in behavior for a path monitoring failure on a new session versus an established session.

Behavior of a session on a monitoring failure	If the rule stays enabled when the monitored IP address is unreachable	If rule is disabled when the monitored IP address is unreachable
For an established session	**wait-recover**—Continue to use egress interface specified in the PBF rule	**wait-recover**—Continue to use egress interface specified in the PBF rule
	fail-over—Use path determined by routing table (no PBF)	**fail-over**—Use path determined by routing table (no PBF)
For a new session	**wait-recover**—Use path determined by routing table (no PBF)	**wait-recover**—Check the remaining PBF rules. If no match, use the routing table
	fail-over—Use path determined by routing table (no PBF)	**fail-over**—Check the remaining PBF rules. If no match, use the routing table

Service Versus Applications in PBF

PBF rules are applied either on the first packet (SYN) or the first response to the first packet (SYN/ACK). This means that a PBF rule may be applied before the firewall has enough information to determine the application. Therefore, application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application.

However, if you specify an application in a PBF rule, the firewall performs *App-ID caching*. When an application passes through the firewall for the first time, the firewall does not have enough information to identify the application and therefore cannot enforce the PBF rule. As more packets arrive, the firewall determines the application and creates an entry in the App-ID cache and retains this App-ID for the session.When a new session is created with the same destination IP address, destination port, and protocol ID, the firewall could identify the application as the same from the initial session (based on the App-ID cache) and apply the PBF rule. Therefore, a session that is not an exact match and is not the same application, can be forwarded based on the PBF rule.

Further, applications have dependencies and the identity of the application can change as the firewall receives more packets. Because PBF makes a routing decision at the start of a session, the firewall cannot enforce a change in application identity. YouTube, for example, starts as web-browsing but changes to Flash, RTSP, or YouTube based on the different links and videos included on the page. However with PBF, because the firewall identifies the application as web-browsing at the start of the session, the change in application is not recognized thereafter.

You cannot use custom-applications, application-filters or application groups in PBF rules.

© 2018 Palo Alto Networks, Inc.

Create a Policy-Based Forwarding Rule

Use a PBF rule to direct traffic to a specific egress interface on the firewall and override the default path for the traffic.

STEP 1 | Create a PBF rule.

When creating a PBF rule you must specify a name for the rule, a source zone or interface, and an egress interface. All other components are either optional or have a default value provided.



You can specify the source and destination addresses using an IP address, an address object, or a FQDN. For the next hop, however, you must specify an IP address.

1. Select **Policies** > **Policy Based Forwarding** and click **Add**.
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, select the following:
   1. Select the **Type**—**Zone** or **Interface**— to which the forwarding policy will be applied, and the relevant zone or interface. If you want to enforce symmetric return, you must select a source interface.

      PBF is only supported on Layer 3 interfaces; loopback interfaces do not support PBF.
   2. (Optional) Specify the **Source Address** to which PBF will apply. For example, a specific IP address or subnet IP address from which you want to forward traffic to the interface or zone specified in this rule.

      Use the Negate option to exclude a one or more source IP addresses from the PBF rule. For example, if your PBF rule directs all traffic from the specified zone to the internet, Negate allows you to exclude internal IP addresses from the PBF rule. (You can also use Negate to exclude destination IP addresses you specify in substep d.)

      The evaluation order is top down. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated.
   3. (Optional) **Add** and select the **Source User** or groups of users to whom the policy applies.
4. In the **Destination/Application/Service** tab, select the following:
   1. **Destination Address**. By default the rule applies to **Any** IP address. Use the **Negate** option to exclude one or more destination IP addresses from the PBF rule.
   2. Select the Application(s) or Service(s) that you want to control using PBF.

      Application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application. For more details, see Service Versus Applications in PBF.

STEP 2 | Specify how to forward traffic that matches the rule.



If you are configuring PBF in a multi-VSYS environment, you must create separate PBF rules for each virtual system (and create the appropriate Security policy rules to enable the traffic).

1. In the **Forwarding** tab, select the following:
   1. Set the **Action.** The options are as follows:

© 2018 Palo Alto Networks, Inc.

- **Forward**—Directs the packet to a specific **Egress Interface**. Enter the **Next Hop** IP address for the packet (you cannot use a domain name for the next hop).
- **Forward To VSYS**—(On a firewall enabled for multiple virtual systems) Select the virtual system to which to forward the packet.
- **Discard**—Drop the packet.
- **No PBF**—Exclude the packets that match the criteria for source/destination/application/ service defined in the rule. Matching packets use the route table instead of PBF; the firewall uses the route table to exclude the matched traffic from the redirected port.

  To trigger the specified action at a daily, weekly or non-recurring frequency, create and attach a Schedule.
- (Optional) Enable Monitoring to verify connectivity to a target IP address or to the next hop IP address. Select **Monitor** and attach a monitoring **Profile** (default or custom) that specifies the action when the IP address is unreachable.

2. (Required for asymmetric routing environments; otherwise, optional) Select **Enforce Symmetric Return** and enter one or more IP addresses in the **Next Hop Address List** (you cannot use an FQDN as the next hop). You can add up to 8 next-hop IP addresses; tunnel and PPoE interfaces are not available as a next-hop IP address.

   Enabling symmetric return ensures that return traffic (say, from the Trust zone on the LAN to the internet) is forwarded out through the same interface through which traffic ingresses from the internet.

STEP 3 | Save the policies to the running configuration on the firewall.

Click **Commit**. The PBF rule is in effect.

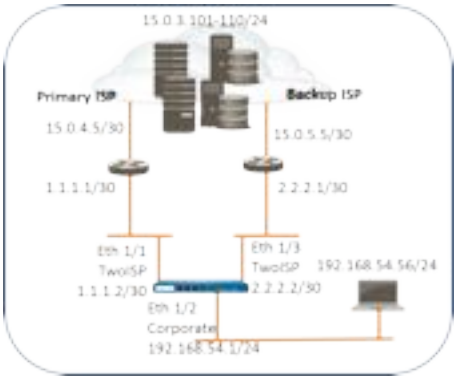


Use Case: PBF for Outbound Access with Dual ISPs

In this use case, the branch office has a dual ISP configuration and implements PBF for redundant internet access. The backup ISP is the default route for traffic from the client to the web servers. In order to enable redundant internet access without using an internetwork protocol such as BGP, we use PBF with destination interface-based source NAT and static routes, and configure the firewall as follows:

- Enable a PBF rule that routes traffic through the primary ISP, and attach a monitoring profile to the rule. The monitoring profile triggers the firewall to use the default route through the backup ISP when the primary ISP is unavailable.
- Define Source NAT rules for both the primary and backup ISP that instruct the firewall to use the source IP address associated with the egress interface for the corresponding ISP. This ensures that the outbound traffic has the correct source IP address.
- Add a static route to the backup ISP, so that when the primary ISP is unavailable, the default route comes into effect and the traffic is directed through the backup ISP.

© 2018 Palo Alto Networks, Inc.



STEP 1 | Configure the ingress and the egress interfaces on the firewall.

Egress interfaces can be in the same zone.

1. Select **Network** > **Interfaces** and then select the interface you want to configure, for example, Ethernet1/1 and Ethernet1/3.

   The interface configuration on the firewall used in this example is as follows:

   - Ethernet 1/1 connected to the primary ISP:
     - Zone: TwoISP
     - IP Address: 1.1.1.2/30
     - Virtual Router: Default
   - Ethernet 1/3 connected to the backup ISP:
     - Zone: TwoISP
     - IP Address: 2.2.2.2/30
     - Virtual Router: Default
   - Ethernet 1/2 is the ingress interface, used by the network clients to connect to the internet:
     - Zone: Corporate
     - IP Address: 192.168. 54.1/24
     - Virtual Router: Default
2. To save the interface configuration, click **OK**.

STEP 2 | On the virtual router, add a static route to the backup ISP.

1. Select **Network** > **Virtual Router** and then select the **default** link to open the Virtual Router dialog.
2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and specify the **Destination** IP address for which you are defining the static route. In this example, we use 0.0.0.0/0 for all traffic.
3. Select the **IP Address** radio button and set the **Next Hop** IP address for your router that connects to the backup internet gateway (you cannot use a domain name for the next hop). In this example, 2.2.2.1.
4. Specify a cost metric for the route. In this example, we use 10.



© 2018 Palo Alto Networks, Inc.

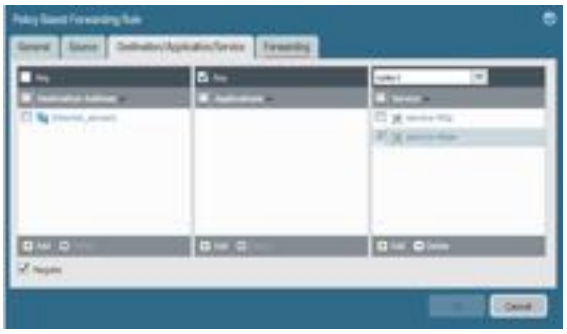
5. Click **OK** twice to save the virtual router configuration.

STEP 3 | Create a PBF rule that directs traffic to the interface that is connected to the primary ISP.

Make sure to exclude traffic destined to internal servers/IP addresses from PBF. Define a negate rule so that traffic destined to internal IP addresses is not routed through the egress interface defined in the PBF rule.

1. Select **Policies** > **Policy Based Forwarding** and click **Add**.
2. Give the rule a descriptive **Name** in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Corporate.
4. In the **Destination/Application/Service** tab, set the following:
    1. In the Destination Address section, **Add** the IP addresses or address range for servers on the internal network or create an address object for your internal servers. Select **Negate** to exclude the IP addresses or address object listed above from using this rule.
    2. In the Service section, **Add** the **service-http** and **service-https** services to allow HTTP and HTTPS traffic to use the default ports. For all other traffic that is allowed by security policy, the default route will be used.

    To forward all traffic using PBF, set the Service to Any.



STEP 4 | Specify where to forward traffic.

1. In the **Forwarding** tab, specify the interface to which you want to forward traffic and enable path monitoring.
2. To forward traffic, set the **Action** to **Forward**, and select the **Egress Interface** and specify the **Next Hop**. In this example, the egress interface is ethernet1/1, and the next hop IP address is 1.1.1.1 (you cannot use a FQDN for the next hop).



3. Enable **Monitor** and attach the default monitoring profile to trigger a failover to the backup ISP. In this example, we do not specify a target IP address to monitor. The firewall will monitor the next hop IP address; if this IP address is unreachable, the firewall will direct traffic to the default route specified on the virtual router.

© 2018 Palo Alto Networks, Inc.

4. (Required if you have asymmetric routes) Select **Enforce Symmetric Return** to ensure that return traffic from the Corporate zone to the internet is forwarded out on the same interface through which traffic ingressed from the internet.
5. NAT ensures that the traffic from the internet is returned to the correct interface/IP address on the firewall.
6. Click **OK** to save the changes.



STEP 5 | Create NAT rules based on the egress interface and ISP. These rules ensure that the correct source IP address is used for outbound connections.

1. Select **Policies** > **NAT** and click **Add**.
2. In this example, the NAT rule we create for each ISP is as follows:

   NAT for Primary ISP

   In the **Original Packet** tab,

   Source Zone: Corporate

   Destination Zone: TwoISP

   In the **Translated Packet** tab, under Source Address Translation

   Translation Type: Dynamic IP and Port

   Address Type: Interface Address

   Interface: ethernet1/1

   IP Address: 1.1.1.2/30

   NAT for Backup ISP

   In the **Original Packet** tab,

   Source Zone: Corporate

   Destination Zone: TwoISP

   In the **Translated Packet** tab, under Source Address Translation

   Translation Type: Dynamic IP and Port

   Address Type: Interface Address

   Interface: ethernet1/2

   IP Address: 2.2.2.2/30



STEP 6 | Create security policy to allow outbound access to the internet.

To safely enable applications, create a simple rule that allows access to the internet and attach the security profiles available on the firewall.

1. Select **Policies** > **Security** and click **Add**.

© 2018 Palo Alto Networks, Inc.

2. Give the rule a descriptive **Name** in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Corporate.
4. In the **Destination** tab, Set the **Destination Zone** to TwoISP.
5. In the **Service/ URL Category** tab, leave the default **application-default**.
6. In the **Actions** tab, complete these tasks:
    1. Set the **Action Setting** to **Allow**.
    2. Attach the default profiles for Antivirus, Anti-Spyware, Vulnerability Protection and URL Filtering, under **Profile Setting.**
7. Under **Options**, verify that logging is enabled at the end of a session. Only traffic that matches a security rule is logged.



STEP 7 | Save the policies to the running configuration on the firewall.

Click **Commit**.

STEP 8 | Verify that the PBF rule is active and that the primary ISP is used for internet access.

1. Launch a web browser and access a web server. On the firewall check the traffic log for web-browsing activity.
2. From a client on the network, use the ping utility to verify connectivity to a web server on the internet. and check the traffic log on the firewall.

```
C:\Users\pm-user1>ping 198.51.100.6
Pinging 198.51.100.6 with 32 bytes of data:
Reply from 198.51.100.6: bytes=32 time=34ms TTL=117
Reply from 198.51.100.6: bytes=32 time=13ms TTL=117
Reply from 198.51.100.6: bytes=32 time=25ms TTL=117
Reply from 198.51.100.6: bytes=32 time=3ms TTL=117
Ping statistics for 198.51.100.6:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 3ms, Maximum = 34ms, Average = 18ms
```



3. To confirm that the PBF rule is active, use the following CLI command:

```
admin@PA-NGFW> show pbf rule all
Rule        ID     Rule State Action   Egress IF/VSYS  NextHop
========== === ========== ====== ==============
Use ISP-Pr 1 Active     Forward ethernet1/1 1.1.1.1
```

STEP 9 | Verify that the failover to the backup ISP occurs and that the Source NAT is correctly applied.

1. Unplug the connection to the primary ISP.
2. Confirm that the PBF rule is inactive with the following CLI command:

```
admin@PA-NGFW> show pbf rule all
Rule        ID     Rule State Action   Egress IF/VSYS  NextHop
========== === ========== ====== ============== ===
Use ISP-Pr 1 Disabled Forward ethernet1/1     1.1.1.1
```


© 2018 Palo Alto Networks, Inc.

3. Access a web server, and check the traffic log to verify that traffic is being forwarded through the backup ISP.



4. View the session details to confirm that the NAT rule is working properly.

```
admin@PA-NGFW> show session all
--------------------------------------------------------------
ID Application    State   Type Flag Src[Sport]/Zone/Proto (translated
 IP[Port]) Vsys Dst[Dport]/Zone (translated IP[Port])
--------------------------------------------------------------
87212 ssl ACTIVE  FLOW  NS   192.168.54.56[53236]/Corporate/6
 (2.2.2.2[12896]) vsys1 204.79.197.200[443]/TwoISP (204.79.197.200[443])
```

5. Obtain the session identification number from the output and view the session details.

The PBF rule is not used and hence is not listed in the output.

```
admin@PA-NGFW> show session id 87212
Session           87212
c2s flow:
                source:      192.168.54.56 [Corporate]
                dst:         204.79.197.200
                proto:       6
                sport:       53236           dport:      443
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
s2c flow:
                source:      204.79.197.200 [TwoISP]
                dst:         2.2.2.2
                proto:       6
                sport:       443             dport:      12896
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
start time                    : Wed Nov5 11:16:10 2014
        timeout                         : 1800 sec
        time to live                    : 1757 sec
        total byte count(c2s)           : 1918
        total byte count(s2c)           : 4333
        layer7 packet count(c2s)        : 10
        layer7 packet count(s2c)        : 7
        vsys                            : vsys1
        application                     : ssl
        rule                            : Corp2ISP
        session to be logged at end     : True
        session in session ager         : True
        session synced from HA peer     : False
        address/port translation        : source
       nat-rule          : NAT-Backup ISP(vsys1)
        layer7 processing               : enabled
        URL filtering enabled           : True
        URL category                    : search-engines
        session via syn-cookies         : False
```


© 2018 Palo Alto Networks, Inc.

```
        session terminated on host    : False
        session traverses tunnel      : False
        captive portal session        : False
        ingress interface       : ethernet1/2
    egress interface       : ethernet1/3
        session QoS rule              : N/A (class 4)
```

Virtual Systems

This topic describes virtual systems, their benefits, typical use cases, and how to configure them. It also provides links to other topics where virtual systems are documented as they function with other features.

> Virtual Systems Overview
> Communication Between Virtual Systems
> Shared Gateway
> Configure Virtual Systems
> Configure Inter-Virtual System Communication within the Firewall
> Configure a Shared Gateway
> Customize Service Routes for a Virtual System
> Virtual System Functionality with Other Features

© 2018 Palo Alto Networks, Inc.

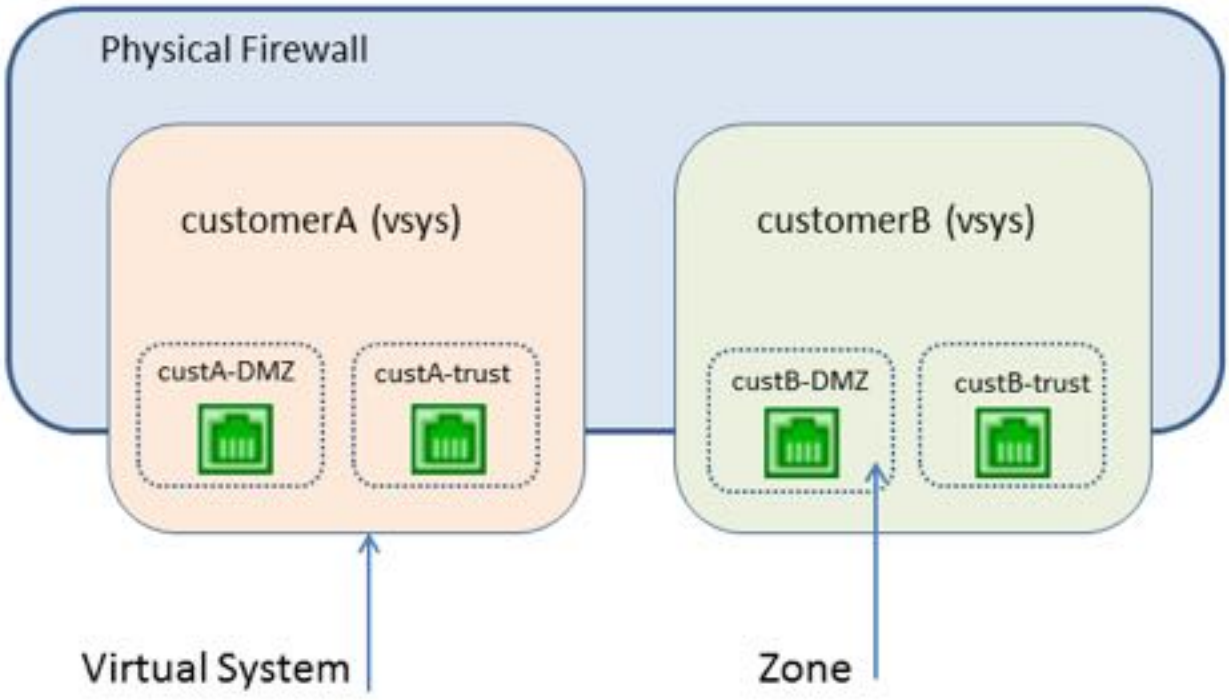
Virtual Systems Overview

Virtual systems are separate, logical firewall instances within a single physical Palo Alto Networks firewall. Rather than using multiple firewalls, managed service providers and enterprises can use a single pair of firewalls (for high availability) and enable virtual systems on them. Each virtual system (vsys) is an independent, separately-managed firewall with its traffic kept separate from the traffic of other virtual systems.

- Virtual System Components and Segmentation
- Benefits of Virtual Systems
- Use Cases for Virtual Systems
- Platform Support and Licensing for Virtual Systems
- Administrative Roles for Virtual Systems
- Shared Objects for Virtual Systems

Virtual System Components and Segmentation

A virtual system is an object that creates an administrative boundary, as shown in the following figure.



A virtual system consists of a set of physical and logical interfaces and subinterfaces (including VLANs and virtual wires), virtual routers, and security zones. You choose the deployment mode(s) (any combination of virtual wire, Layer 2, or Layer 3) of each virtual system. By using virtual systems, you can segment any of the following:

- Administrative access
- The management of all policies (Security, NAT, QoS, Policy-based Forwarding, Decryption, Application Override, Authentication, and DoS protection)
- All objects (such as address objects, application groups and filters, dynamic block lists, security profiles, decryption profiles, custom objects, etc.)
- User-ID

© 2018 Palo Alto Networks, Inc.

- Certificate management
- Server profiles
- Logging, reporting, and visibility functions

Virtual systems affect the security functions of the firewall, but virtual systems alone do not affect networking functions such as static and dynamic routing. You can segment routing for each virtual system by creating one or more virtual routers for each virtual system, as in the following use cases:

- If you have virtual systems for departments of one organization, and the network traffic for all of the departments is within a common network, you can create a single virtual router for multiple virtual systems.
- If you want routing segmentation and each virtual system's traffic must be isolated from other virtual systems, you can create one or more virtual routers for each virtual system.

Benefits of Virtual Systems

Virtual systems provide the same basic functions as a physical firewall, along with additional benefits:

- **Segmented administration**—Different organizations (or customers or business units) can control (and monitor) a separate firewall instance, so that they have control over their own traffic without interfering with the traffic or policies of another firewall instance on the same physical firewall.
- **Scalability**—After the physical firewall is configured, adding or removing customers or business units can be done efficiently. An ISP, managed security service provider, or enterprise can provide different security services to each customer.
- **Reduced capital and operational expenses**—Virtual systems eliminate the need to have multiple physical firewalls at one location because virtual systems co-exist on one firewall. By not having to purchase multiple firewalls, an organization can save on the hardware expense, electric bills, and rack space, and can reduce maintenance and management expenses.

Use Cases for Virtual Systems

There are many ways to use virtual systems in a network. One common use case is for an ISP or a managed security service provider (MSSP) to deliver services to multiple customers with a single firewall. Customers can choose from a wide array of services that can be enabled or disabled easily. The firewall's role-based administration allows the ISP or MSSP to control each customer's access to functionality (such as logging and reporting) while hiding or offering read-only capabilities for other functions.

Another common use case is within a large enterprise that requires different firewall instances because of different technical or confidentiality requirements among multiple departments. Like the above case, different groups can have different levels of access while IT manages the firewall itself. Services can be tracked and/or billed back to departments to thereby make separate financial accountability possible within an organization.

Platform Support and Licensing for Virtual Systems

Virtual systems are supported on the PA-3000 Series, PA-5000 Series, PA-5200 Series, and PA-7000 Series firewalls. Each firewall series supports a base number of virtual systems; the number varies by platform. A Virtual Systems license is required to support multiple virtual systems on the PA-3000 Series firewalls, and to create more than the base number of virtual systems supported on a platform.

For license information, see Activate Licenses and Subscriptions. For the base and maximum number of virtual systems supported, see Compare Firewalls tool.

Multiple virtual systems are not supported on the PA-200, PA-220, PA-500, PA-800 Series, or VM-Series firewalls.

© 2018 Palo Alto Networks, Inc.

Administrative Roles for Virtual Systems

A **superuser** administrator can create virtual systems and add a **Device Administrator**, **vsysadmin**, or **vsysreader**. A **Device Administrator** can access all virtual systems, but cannot add administrators. The two types of virtual system administrative roles are:

- **vsysadmin**—Grants full access to a virtual system.
- **vsysreader**—Grants read-only access to a virtual system.

A virtual system administrator can view logs of only the virtual systems assigned to that administrator. Someone with **superuser** or **Device Admin** permission can view all of the logs or select a virtual system to view.

Persons with **vsysadmin** permission can commit configurations for only the virtual systems assigned to them.

Shared Objects for Virtual Systems

If your administrator account extends to multiple virtual systems, you can choose to configure objects (such as an address object) and policies for a specific virtual system or as shared objects, which apply to all of the virtual systems on the firewall. If you try to create a shared object with the same name and type as an existing object in a virtual system, the virtual system object is used.

© 2018 Palo Alto Networks, Inc.

Communication Between Virtual Systems

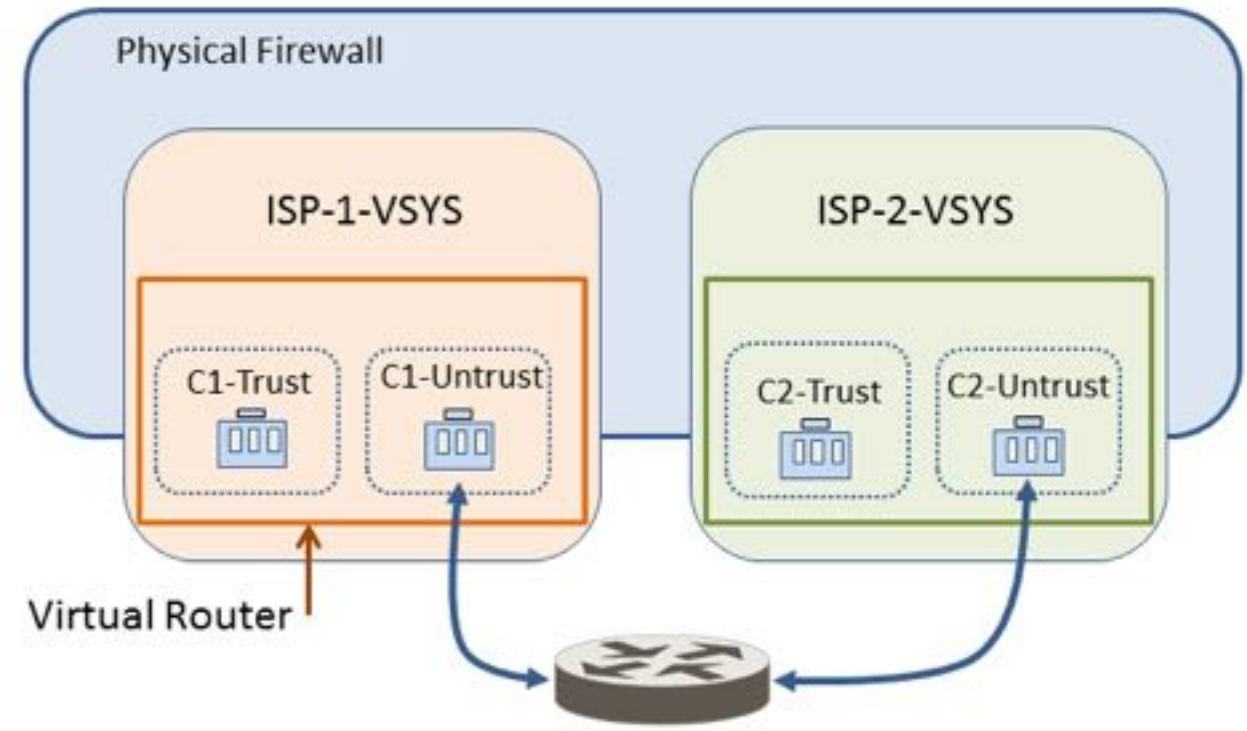
There are two typical scenarios where communication between virtual systems (inter-vsys traffic) is desirable. In a multi-tenancy environment, communication between virtual systems can occur by having traffic leave the firewall, go through the Internet, and re-enter the firewall. In a single organization environment, communication between virtual systems can remain within the firewall. This section discusses both scenarios.

- Inter-VSYS Traffic That Must Leave the Firewall on page 1048
- Inter-VSYS Traffic That Remains Within the Firewall on page 1048
- Inter-VSYS Communication Uses Two Sessions on page 1050

Inter-VSYS Traffic That Must Leave the Firewall

An ISP that has multiple customers on a firewall (known as multi-tenancy) can use a virtual system for each customer, and thereby give each customer control over its virtual system configuration. The ISP grants **vsysadmin** permission to customers. Each customer's traffic and management are isolated from the others. Each virtual system must be configured with its own IP address and one or more virtual routers in order to manage traffic and its own connection to the Internet.

If the virtual systems need to communicate with each other, that traffic goes out the firewall to another Layer 3 routing device and back to the firewall, even though the virtual systems exist on the same physical firewall, as shown in the following figure.



Inter-VSYS Traffic That Remains Within the Firewall

Unlike the preceding multi-tenancy scenario, virtual systems on a firewall can be under the control of a single organization. The organization wants to both isolate traffic between virtual systems and allow communications between virtual systems. This common use case arises when the organization wants to

© 2018 Palo Alto Networks, Inc.

provide departmental separation and still have the departments be able to communicate with each other or connect to the same network(s). In this scenario, the inter-vsys traffic remains within the firewall, as described in the following topics:

- External Zone on page 1049
- External Zones and Security Policies For Traffic Within a Firewall on page 1049

External Zone

The communication desired in the use case above is achieved by configuring security policies that point to or from an *external* zone. An external zone is a security object that is associated with a specific virtual system that it can reach; the zone is external to the virtual system. A virtual system can have only one external zone, regardless of how many security zones the virtual system has within it. External zones are required to allow traffic between zones in different virtual systems, without the traffic leaving the firewall.

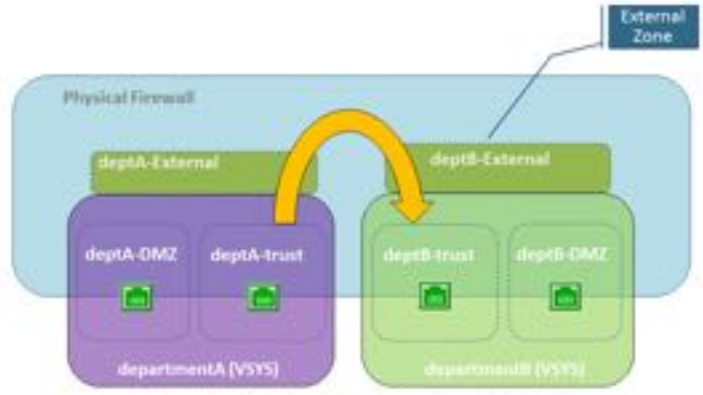
The virtual system administrator configures the security policies needed to allow traffic between two virtual systems. Unlike security zones, an external zone is not associated with an interface; it is associated with a virtual system. The security policy allows or denies traffic between the security (internal) zone and the external zone.

Because external zones do not have interfaces or IP addresses associated with them, some zone protection profiles are not supported on external zones.

Remember that each virtual system is a separate instance of a firewall, which means that each packet moving between virtual systems is inspected for security policy and App-ID evaluation.

External Zones and Security Policies For Traffic Within a Firewall

In the following example, an enterprise has two separate administrative groups: the departmentA and departmentB virtual systems. The following figure shows the external zone associated with each virtual system, and traffic flowing from one trust zone, out an external zone, into an external zone of another virtual system, and into its trust zone.



To create external zones, the firewall administrator must configure the virtual systems so that they are *visible* to each other. External zones do not have security policies between them because their virtual systems are visible to each other.

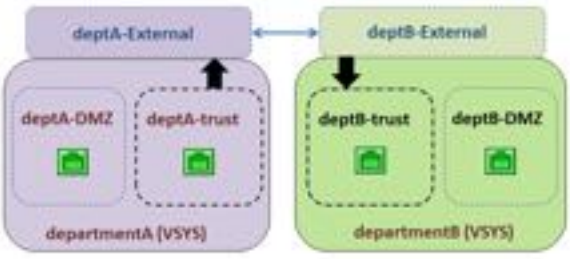
To communicate between virtual systems, the ingress and egress interfaces on the firewall are either assigned to a single virtual router or else they are connected using inter-virtual router static routes. The simpler of these two approaches is to assign all virtual systems that must communicate with each other to a single virtual router.

There might be a reason that the virtual systems need to have their own virtual router, for example, if the virtual systems use overlapping IP address ranges. Traffic can be routed between the virtual systems, but each virtual router must have static routes that point to the other virtual router(s) as the next hop.

© 2018 Palo Alto Networks, Inc.

Referring to the scenario in the figure above, we have an enterprise with two administrative groups: departmentA and departmentB. The departmentA group manages the local network and the DMZ resources. The departmentB group manages traffic in and out of the sales segment of the network. All traffic is on a local network, so a single virtual router is used. There are two external zones configured for communication between the two virtual systems. The departmentA virtual system has three zones used in security policies: deptA-DMZ, deptA-trust, and deptA-External. The departmentB virtual system also has three zones: deptB-DMZ, deptB-trust, and deptB-External. Both groups can control the traffic passing through their virtual systems.

In order to allow traffic from deptA-trust to deptB-trust, two security policies are required. In the following figure, the two vertical arrows indicate where the security policies (described below the figure) are controlling traffic.



- Security Policy 1: In the preceding figure, traffic is destined for the deptB-trust zone. Traffic leaves the deptA-trust zone and goes to the deptA-External zone. A security policy must allow traffic from the source zone (deptA-trust) to the destination zone (deptA-External). A virtual system allows any policy type to be used for this traffic, including NAT.

  No policy is needed between external zones because traffic sent to an external zone appears in and has automatic access to the other external zones that are visible to the original external zone.
- Security Policy 2: In the preceding figure, the traffic from deptB-External is still destined to the deptB-trust zone, and a security policy must be configured to allow it. The policy must allow traffic from the source zone (deptB-External) to the destination zone (deptB-trust).

The departmentB virtual system could be configured to block traffic from the departmentA virtual system, and vice versa. Like traffic from any other zone, traffic from external zones must be explicitly allowed by policy to reach other zones in a virtual system.

In addition to external zones being required for inter-virtual system traffic that does not leave the firewall, external zones are also required if you configure a Shared Gateway*, in which case the traffic is intended to leave the firewall.*

Inter-VSYS Communication Uses Two Sessions

It is helpful to understand that communication between two virtual systems uses two sessions, unlike the one session used for a single virtual system. Let's compare the scenarios.

Scenario 1—Vsys1 has two zones: trust1 and untrust1. A host in the trust1 zone initiates traffic when it needs to communicate with a device in the untrust1 zone. The host sends traffic to the firewall, and the firewall creates a new session for source zone trust1 to destination zone untrust1. Only one session is needed for this traffic.

Scenario 2—A host from vsys1 needs to access a server on vsys2. A host in the trust1 zone initiates traffic to the firewall, and the firewall creates the first session: source zone trust1 to destination zone untrust1. Traffic is routed to vsys2, either internally or externally. Then the firewall creates a second session: source zone untrust2 to destination zone trust2. Two sessions are needed for this inter-vsys traffic.

© 2018 Palo Alto Networks, Inc.

Shared Gateway

This topic includes the following information about shared gateways:

- External Zones and Shared Gateway on page 1051
- Networking Considerations for a Shared Gateway on page 1052

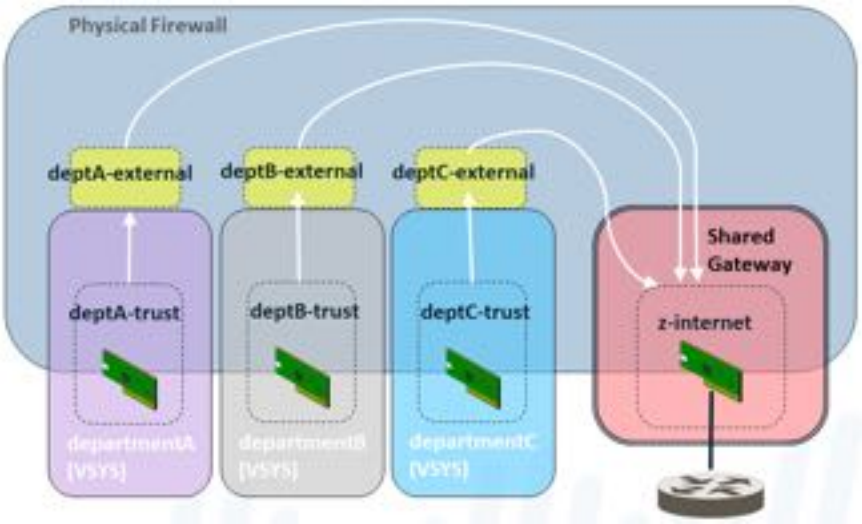
External Zones and Shared Gateway

A shared gateway is an interface that multiple virtual systems share in order to communicate over the Internet. Each virtual system requires an External Zone, which acts as an intermediary, for configuring security policies that allow or deny traffic from the virtual system's internal zone to the shared gateway.

The shared gateway uses a single virtual router to route traffic for all virtual systems. A shared gateway is used in cases when an interface does not need a full administrative boundary around it, or when multiple virtual systems must share a single Internet connection. This second case arises if an ISP provides an organization with only one IP address (interface), but multiple virtual systems need external communication.

Unlike the behavior between virtual systems, security policy and App-ID evaluations are not performed between a virtual system and a shared gateway. That is why using a shared gateway to access the Internet involves less overhead than creating another virtual system to do so.

In the following figure, three customers share a firewall, but there is only one interface accessible to the Internet. Creating another virtual system would add the overhead of App-ID and security policy evaluation for traffic being sent to the interface through the added virtual system. To avoid adding another virtual system, the solution is to configure a shared gateway, as shown in the following diagram.



The shared gateway has one globally-routable IP address used to communicate with the outside world. Interfaces in the virtual systems have IP addresses too, but they can be private, non-routable IP addresses.

You will recall that an administrator must specify whether a virtual system is visible to other virtual systems. Unlike a virtual system, a shared gateway is always visible to all of the virtual systems on the firewall.

A shared gateway ID number appears as **sg<ID** on the web interface. It is recommended that you name your shared gateway with a name that includes its ID number.

When you add objects such as zones or interfaces to a shared gateway, the shared gateway appears as an available virtual system in the vsys drop-down menu.

© 2018 Palo Alto Networks, Inc.

A shared gateway is a limited version of a virtual system; it supports NAT and policy-based forwarding (PBF), but does not support Security, DoS policies, QoS, Decryption, Application Override, or Authentication policies.

Networking Considerations for a Shared Gateway

Keep the following in mind while you are configuring a shared gateway.

- The virtual systems in a shared gateway scenario access the Internet through the shared gateway's physical interface, using a single IP address. If the IP addresses of the virtual systems are not globally routable, configure source NAT to translate those addresses to globally-routable IP addresses.
- A virtual router routes the traffic for all of the virtual systems through the shared gateway.
- The default route for the virtual systems should point to the shared gateway.
- Security policies must be configured for each virtual system to allow the traffic between the internal zone and external zone, which is visible to the shared gateway.
- A firewall administrator should control the virtual router, so that no member of a virtual system can affect the traffic of other virtual systems.
- Within a Palo Alto Networks firewall, a packet may hop from one virtual system to another virtual system or a shared gateway. A packet may not traverse more than two virtual systems or shared gateways. For example, a packet cannot go from one virtual system to a shared gateway to a second virtual system within the firewall.

To save configuration time and effort, consider the following advantages of a shared gateway:

- Rather than configure NAT for multiple virtual systems associated with a shared gateway, you can configure NAT for the shared gateway.
- Rather than configure policy-based routing (PBR) for multiple virtual systems associated with a shared gateway, you can configure PBR for the shared gateway.

© 2018 Palo Alto Networks, Inc.

Configure Virtual Systems

Creating a virtual system requires that you have the following:

- A **superuser** administrative role.
- An interface configured.
- A Virtual Systems license if you are configuring a PA-3000 Series firewall, or if you are creating more than the base number of virtual systems supported on the platform. See Platform Support and Licensing for Virtual Systems.

STEP 1 | Enable virtual systems.

1. Select **Device** > **Setup** > **Management** and edit the **General Settings**.
2. Select the **Multi Virtual System Capability** check box and click **OK**. This action triggers a commit if you approve it.

   Only after enabling virtual systems will the **Device** tab display the **Virtual Systems** and **Shared Gateways** options.

STEP 2 | Create a virtual system.

1. Select **Device** > **Virtual Systems**, click **Add** and enter a virtual system **ID**, which is appended to "vsys" (range is 1-255).

   The default ID is 1, which makes the default virtual system vsys1. This default appears even on platforms that do not support multiple virtual systems.
2. Select **Allow forwarding of decrypted content** if you want to allow the firewall to forward decrypted content to an outside service. For example, you must enable this option for the firewall to be able to send decrypted content to WildFire for analysis.
3. Enter a descriptive **Name** for the virtual system. A maximum of 31 alphanumeric, space, and underscore characters is allowed.

STEP 3 | Assign interfaces to the virtual system.

The virtual routers, virtual wires, or VLANs can either be configured already or you can configure them later, at which point you specify the virtual system associated with each.

1. On the **General** tab, select a **DNS Proxy** object if you want to apply DNS proxy rules to the interface.
2. In the **Interfaces** field, click **Add** to enter the interfaces or subinterfaces to assign to the virtual system. An interface can belong to only one virtual system.
3. Do any of the following, based on the deployment type(s) you need in the virtual system:
   - In the **VLANs** field, click **Add** to enter the VLAN(s) to assign to the vsys.
   - In the **Virtual Wires** field, click **Add** to enter the virtual wire(s) to assign to the vsys.
   - In the **Virtual Routers** field, click **Add** to enter the virtual router(s) to assign to the vsys.
4. In the **Visible Virtual System** field, check all virtual systems that should be made visible to the virtual system being configured. This is required for virtual systems that need to communicate with each other.

   In a multi-tenancy scenario where strict administrative boundaries are required, no virtual systems would be checked.
5. Click **OK**.

© 2018 Palo Alto Networks, Inc.

STEP 4 | (Optional) Limit the resource allocations for sessions, rules, and VPN tunnels allowed for the virtual system. The flexibility of being able to allocate limits per virtual system allows you to effectively control firewall resources.

1. On the **Resource** tab, optionally set limits for a virtual system. Each field displays the valid range of values; there are no default values.
   - **Sessions Limit**

     If you use the show session meter CLI command, it displays the Maximum number of sessions allowed per dataplane, the Current number of sessions being used by the virtual system, and the Throttled number of sessions per virtual system. On a PA-5200 or PA-7000 Series firewall, the Current number of sessions being used can be greater than the Maximum configured for Sessions Limit because there are multiple dataplanes per virtual system. The Sessions Limit you configure on a PA-5200 Series or PA-7000 Series firewall is per dataplane, and will result in a higher maximum per virtual system.
   - **Security Rules**
   - **NAT Rules**
   - **Decryption Rules**
   - **QoS Rules**
   - **Application Override Rules**
   - **Policy Based Forwarding Rules**
   - **Authentication Rules**
   - **DoS Protection Rules**
   - **Site to Site VPN Tunnels**
   - **Concurrent SSL VPN Tunnels**
2. Click **OK**.

STEP 5 | Commit the configuration.

Click **Commit**. The virtual system is now an object accessible from the **Objects** tab.

STEP 6 | Create at least one virtual router for the virtual system in order to make the virtual system capable of networking functions, such as static and dynamic routing.

Alternatively, your virtual system might use a VLAN or a virtual wire, depending on your deployment.

1. Select **Network** > **Virtual Routers** and **Add** a virtual router by **Name**.
2. For **Interfaces**, click **Add** and from the drop-down, select the interfaces that belong to the virtual router.
3. Click **OK**.

STEP 7 | Configure a security zone for each interface in the virtual system.

For at least one interface, create a Layer 3 security zone. See Configure Interfaces and Zones.

STEP 8 | Configure the security policy rules that allow or deny traffic to and from the zones in the virtual system.

See Create a Security Policy Rule.

STEP 9 | Commit the configuration.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.



After creating a virtual system, you can use the CLI to commit a configuration for only a specific virtual system:

commit partial vsys *<vsys-id>*

STEP 10 | (Optional) View the security policies configured for a virtual system.

Open an SSH session to use the CLI. To view the security policies for a virtual system, in operational mode, use the following commands:

set system setting target-vsys *<vsys-id>*

show running security-policy

© 2018 Palo Alto Networks, Inc.

Configure Inter-Virtual System Communication within the Firewall

Perform this task if you have a use case, perhaps within a single enterprise, where you want the virtual systems to be able to communicate with each other within the firewall. Such a scenario is described in Inter-VSYS Traffic That Remains Within the Firewall. This task presumes:

- You completed the task, Configure Virtual Systems.
- When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate with each other to be visible to each other.

STEP 1 | Configure an external zone for each virtual system.

1. Select **Network** > **Zones** and **Add** a new zone by **Name**.
2. For **Location**, select the virtual system for which you are creating an external zone.
3. For **Type**, select **External**.
4. For **Virtual Systems**, click **Add** and enter the virtual system that the external zone can reach.
5. (Optional) Select a **Zone Protection Profile** (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.
6. (Optional) In **Log Setting**, select a log forwarding profile for forwarding zone protection logs to an external system.
7. (Optional) Select **Enable User Identification** to enable User-ID for the external zone.
8. Click **OK**.

STEP 2 | Configure the Security policy rules to allow or deny traffic from the internal zones to the external zone of the virtual system, and vice versa.

- See Create a Security Policy Rule.
- See Inter-VSYS Traffic That Remains Within the Firewall.

STEP 3 | Commit your changes.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Configure a Shared Gateway

Perform this task if you need multiple virtual systems to share an interface (a Shared Gateway) to the Internet. This task presumes:

- You configured an interface with a globally-routable IP address, which will be the shared gateway.
- You completed the prior task, Configure Virtual Systems. For the interface, you chose the external-facing interface with the globally-routable IP address.
- When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate to be visible to each other.

STEP 1 | Configure a Shared Gateway.

1. Select **Device** > **Shared Gateway**, click **Add** and enter an **ID**.
2. Enter a helpful **Name**, preferably including the **ID** of the gateway.
3. In the **DNS Proxy** field, select a DNS proxy object if you want to apply DNS proxy rules to the interface.
4. **Add** an **Interface** that connects to the outside world.
5. Click **OK**.

STEP 2 | Configure the zone for the shared gateway.



When adding objects such as zones or interfaces to a shared gateway, the shared gateway itself will be listed as an available vsys in the VSYS drop-down menu.

1. Select **Network** > **Zones** and **Add** a new zone by **Name**.
2. For **Location**, select the shared gateway for which you are creating a zone.
3. For **Type**, select **Layer3**.
4. (Optional) Select a **Zone Protection Profile** (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.
5. (Optional) In **Log Setting**, select a log forwarding profile for forwarding zone protection logs to an external system.
6. (Optional) Select **Enable User Identification** to enable User-ID for the shared gateway.
7. Click **OK**.

STEP 3 | Commit your changes.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Customize Service Routes for a Virtual System

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 service 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 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.

- Customize Service Routes to Services for Virtual Systems
- Configure a PA-7000 Series Firewall for Logging Per Virtual System
- Configure Administrative Access Per Virtual System or Firewall

Customize Service Routes to Services for Virtual Systems

When you enable Multi Virtual System Capability, any virtual system that does not have specific service routes configured inherits the global service and service route settings for the firewall. You can instead configure a virtual system to use a different service route, as described in the following workflow.

A firewall with multiple virtual systems must have interfaces and subinterfaces with non-overlapping IP addresses. A per-virtual system service route for SNMP traps or for Kerberos is for IPv4 only.

The service route for a service strictly follows how you configured the server profile for the service:

- If you define a server profile (**Device** > **Server Profiles**) for the Shared location, the firewall uses the global service route for that service.
- If you define a server profile for a specific virtual system, the firewall uses the virtual system-specific service route for that service.
- If you define a server profile for a specific virtual system but the virtual system-specific service route for that service is not configured, the firewall uses the global service route for that service.

The firewall supports syslog forwarding on a virtual system basis. When multiple virtual systems on a firewall are connecting to a syslog server using SSL transport, the firewall can generate only one certificate for secure communication. The firewall does not support each virtual system having its own certificate.

STEP 1 | Customize service routes for a virtual system.

1. Select **Device** > **Setup** > **Services** > **Virtual Systems**, and select the virtual system you want to configure.

© 2018 Palo Alto Networks, Inc.

2. Click the **Service Route Configuration** link.
3. Select one:
   - **Inherit Global Service Route Configuration**—Causes the virtual system to inherit the global service route settings relevant to a virtual system. If you choose this option, skip the step to customize.
   - **Customize**—Allows you to specify a source address for each service.
4. If you chose **Customize**, select the **IPv4** or **IPv6** tab, depending on what type of addressing the server offering the service uses. You can specify both IPv4 and IPv6 addresses for a service. Click on a service. (Only services that are relevant to a virtual system are available.)

   To easily use the same source address for multiple services, select the checkbox for the services, click Set Selected Routes, and continue.

   - To limit the drop-down list for Source Address, select a **Source Interface**, then select a Source Address (from that interface) as the service route. Selecting **Any** Source Interface makes all IP addresses on all interfaces for the virtual system available in the Source Address drop-down list from which you select an address. You can select **Inherit Global Setting**.
   - **Source Address** will indicate **Inherited** if you selected **Inherit Global Setting** for the **Source Interface** or it will indicate the source address you selected. If you selected **Any** for **Source Interface**, select an IP address from the drop-down, or enter an IP address (using the IPv4 or IPv6 format that matches the tab you chose) to specify the source address that will be used in packets sent to the external service.
   - If you modify an address object and the IP family type (IPv4/IPv6) changes, a **Commit** is required to update the service route family to use.
5. Click **OK**.
6. Repeat the prior steps to configure source addresses for other external services.
7. Click **OK**.

STEP 2 | Commit your changes.

Click **Commit** and **OK**.

If you are configuring per-virtual system service routes for logging services for a PA-7000 Series firewall, continue to the task Configure a PA-7000 Series Firewall for Logging Per Virtual System.

Configure a PA-7000 Series Firewall for Logging Per Virtual System

For Traffic, HIP Match, Threat, and Wildfire log types, the PA-7000 Series firewall does not use service routes for SNMP Trap, Syslog and email services. Instead, the PA-7000 Series firewall Log Processing Card (LPC) supports virtual system-specific paths from LPC subinterfaces to an on-premise switch to the respective service on a server. For System and Config logs, the PA-7000 Series firewall uses global service routes, and not the LPC.

In other Palo Alto Networks models, the dataplane sends logging service route traffic to the management plane, which sends the traffic to logging servers. In the PA-7000 Series firewall, each LPC has only one interface, and data planes for multiple virtual systems send logging server traffic (types mentioned above) to the PA-7000 Series firewall LPC. The LPC is configured with multiple subinterfaces, over which the platform sends the logging service traffic out to a customer's switch, which can be connected to multiple logging servers.

Each LPC subinterface can be configured with a subinterface name and a dotted subinterface number. The subinterface is assigned to a virtual system, which is configured for logging services. The other service

© 2018 Palo Alto Networks, Inc.

routes on a PA-7000 Series firewall function similarly to service routes on other Palo Alto Networks platforms. For information about the LPC itself, see the PA-7000 Series Hardware Reference Guide.

If you have enabled multi virtual system capability on your PA-7000 Series firewall, you can configure logging for different virtual systems as described in the following workflow.

STEP 1 | Create a Log Card subinterface.

1. Select **Network** > **Interfaces** > **Ethernet** and select the interface that will be the Log Card interface.
2. Enter the **Interface Name**.
3. For **Interface Type**, select **Log Card** from the drop-down.
4. Click **OK**.

STEP 2 | Add a subinterface for each tenant on the LPCs physical interface.

1. Highlight the Ethernet interface that is a Log Card interface type and click **Add Subinterface**.
2. For **Interface Name**, after the period, enter the subinterface assigned to the tenant's virtual system.
3. For **Tag**, enter a VLAN tag value.

   Make the tag the same as the subinterface number for ease of use, but it could be a different number.
4. (Optional) Enter a **Comment**.
5. On the **Config** tab, in the **Assign Interface to Virtual System** field, select the virtual system to which the LPC subinterface is assigned (from the drop-down). Alternatively, you can click **Virtual Systems** to add a new virtual system.
6. Click **OK**.

STEP 3 | Enter the addresses assigned to the subinterface, and configure the default gateway.

1. Select the **Log Card Forwarding** tab, and do one or both of the following:
   - For the IPv4 section, enter the **IP Address** and **Netmask** assigned to the subinterface. Enter the **Default Gateway** (the next hop where packets will be sent that have no known next hop address in the Routing Information Base [RIB]).
   - For the IPv6 section, enter the **IPv6 Address** assigned to the subinterface. Enter the **IPv6 Default Gateway**.
2. Click **OK**.

STEP 4 | Commit your changes.

Click **OK** and **Commit**.

STEP 5 | If you haven't already done so, configure the remaining service routes for the virtual system.

Customize Service Routes for a Virtual System.

Configure Administrative Access Per Virtual System or Firewall

If you have a superuser administrative account, you can create and configure granular permissions for a vsysadmin or device admin role.

STEP 1 | Create an Admin Role Profile that grants or disables permission to an Administrator to configure or read-only various areas of the web interface.

1. Select **Device** > **Admin Roles** and **Add** an **Admin Role Profile**.
2. Enter a **Name** and optional **Description** of the profile.
3. For **Role**, specify which level of control the profile affects:

© 2018 Palo Alto Networks, Inc.

- **Device**—The profile allows the management of the global settings and any virtual systems.
- **Virtual System**—The profile allows the management of only the virtual system(s) assigned to the administrator(s) who have this profile. (The administrator will be able to access **Device** > **Setup** > **Services** > **Virtual Systems**, but not the **Global** tab.)

4. On the **Web UI** tab for the Admin Role Profile, scroll down to **Device**, and leave the green check mark (Enable).
   - Under **Device**, enable **Setup**. Under **Setup**, enable the areas to which this profile will grant configuration permission to the administrator, as shown below. (The Read Only lock icon appears in the Enable/Disable rotation if Read Only is allowed for that setting.)
     - **Management**—Allows an admin with this profile to configure settings on the **Management** tab.
     - **Operations**—Allows an admin with this profile to configure settings on the **Operations** tab.
     - **Services**—Allows an admin with this profile to configure settings on the **Services** tab. An admin must have **Services** enabled in order to access the **Device** > **Setup Services** > **Virtual Systems** tab. If the **Role** was specified as **Virtual System** in the prior step, **Services** is the only setting that can be enabled under **Device** > **Setup**.
     - **Content-ID**—Allows an admin with this profile to configure settings on the **Content-ID** tab.
     - **WildFire**—Allows an admin with this profile to configure settings on the **WildFire** tab.
     - **Session**—Allows an admin with this profile to configure settings on the **Session** tab.
     - **HSM**—Allows an admin with this profile to configure settings on the **HSM** tab.
5. Click **OK**.
6. (Optional) Repeat the entire step to create another Admin Role profile with different permissions, as necessary.

STEP 2 | Apply the Admin role profile to an administrator.

1. Select **Device** > **Administrators**, click **Add** and enter the **Name** to add an Administrator.
2. (Optional) Select an **Authentication Profile**.
3. (Optional) Select **Use only client certificate authentication (Web)** to have bi-directional authentication; to get the server to authenticate the client.
4. Enter a **Password** and **Confirm Password**.
5. (Optional) Select **Use Public Key Authentication (SSH)** if you want to use a much stronger, key-based authentication method using an SSH public key rather than just a password.
6. For **Administrator Type**, select **Role Based**.
7. For **Profile**, select the profile that you just created.
8. (Optional) Select a **Password Profile**.
9. Click **OK**.

STEP 3 | Commit the configuration.

Click **Commit**.

© 2018 Palo Alto Networks, Inc.

Virtual System Functionality with Other Features

Many firewall features and functionality are capable of being configured, viewed, logged, or reported per virtual system. Therefore, virtual systems are mentioned in other relevant locations in the documentation and that information is not repeated here. Some of the specific chapters are the following:

- If you are configuring Active/Passive HA, the two firewalls must have the same virtual system capability (single or multiple virtual system capability). See High Availability on page 249.
- To configure QoS for virtual systems, see Configure QoS for a Virtual System on page 714.
- For information about configuring a firewall with virtual systems in a virtual wire deployment that uses subinterfaces (and VLAN tags), see Virtual Wire Interfaces on page 799.

Zone Protection and DoS Protection

Attacks against your network can originate externally or internally. Because different parts of a network perform functions that require different types and levels of protection, a global security policy or port-based security is not granular enough to properly secure each part of the network.

The solution is to segment the network into functional and organizational zones to reduce the network's attack surface (the portion of the network and its traffic exposed to potential external and internal attackers). You protect each zone by using zone protection to protect zone borders and denial-of-service (DoS) protection to defend the endpoints and resources in each security zone.

> Network Segmentation Using Zones on page 1065
> How Do Zones Protect the Network? on page 1066
> Zone Defense on page 1067
> Configure Zone Protection to Increase Network Security on page 1074
> DoS Protection Against Flooding of New Sessions on page 1081

© 2018 Palo Alto Networks, Inc.

Network Segmentation Using Zones

The larger the network, the more difficult it is to protect. A large, unsegmented network presents a large attack surface with more weaknesses and vulnerabilities. Because traffic and applications have access to the entire network, once an attacker gains entry to a network, the attacker can move laterally through the network to access critical data. A large network is also more difficult to monitor and control. Segmenting the network limits an attacker's ability to move through the network by preventing lateral movement between zones.

A security zone is a group of one or more physical or virtual firewall interfaces and the network segments connected to the zone's interfaces. You control protection for each zone individually so that each zone receives the specific protections it needs. For example, a zone for the finance department may not need to allow all of the applications that a zone for IT allows.

To fully protect your network, all traffic must flow through the firewall. Configure Interfaces and Zones to create separate zones for different functional areas such as the internet gateway, sensitive data storage, and business applications, and for different organizational groups such as finance, IT, marketing, and engineering. Wherever there is a logical division of functionality, application usage, or user access privileges, you can create a separate zone to isolate and protect the area and apply the appropriate security policy rules to prevent unnecessary access to data and applications that only one or some groups need to access. The more granular the zones, the greater the visibility and control you have over network traffic. Dividing your network into zones helps to create a Zero Trust architecture that executes a security philosophy of trusting no users, devices, applications, or packets, and verifying everything. The end goal is to create a network that allows access only to the users, devices, and applications that have legitimate business needs, and to deny all other traffic.

How to appropriately restrict and permit access to zones depends on the network environment. For example, environments such as semiconductor manufacturing floors or robotic assembly plants, where the workstations control sensitive manufacturing equipment, or highly restricted access areas, may require physical segmentation that permits no access from outside devices (no mobile device access).

In environments where users can access the network with mobile devices, enabling User-ID and App-ID in conjunction with segmenting the network into zones ensures that users receive the appropriate access privileges regardless of where they access the network, because access privileges are tied to a user or a user group instead of to a device in one particular zone.

The protection requirements for different functional areas and groups may also differ. For example, a zone that handles a large amount of traffic may require different flood protection thresholds than a zone that normally handles less traffic. The ability to define the appropriate protection for each zone is another reason to segment the network. What appropriate protection is depends on your network architecture, what you want to protect, and what traffic you want to permit and deny.

© 2018 Palo Alto Networks, Inc.

How Do Zones Protect the Network?

Zones not only protect your network by segmenting it into smaller, more easily controlled areas, zones also protect the network because you can control access to zones and traffic movement between zones.

Zones prevent uncontrolled traffic from flowing through the firewall interfaces into your network because firewall interfaces can't process traffic until you assign them to zones. The firewall applies zone protection on ingress interfaces, where traffic enters the firewall in the direction of flow from the originating client to the responding server (c2s), to filter traffic before it enters a zone.

The firewall interface type and the zone type (Tap, virtual wire, L2, L3, Tunnel, or External) must match, which helps to protect the network against admitting traffic that doesn't belong in a zone. For example, you can assign an L2 interface to an L2 zone or an L3 interface to an L3 zone, but you can't assign an L2 interface to an L3 zone.

In addition, a firewall interface can belong to one zone only. Traffic destined for different zones can't use the same interface, which helps to prevent inappropriate traffic from entering a zone and enables you to configure the protection appropriate for each individual zone. You can connect more than one firewall interface to a zone to increase bandwidth, but each interface can connect to only one zone.

After the firewall admits traffic to a zone, traffic flows freely within that zone and is not logged. The smaller you make each zone, the greater the control you have over the traffic that accesses each zone, and the more difficult it is for malware to move laterally across the network between zones. Traffic can't flow between zones unless a security policy rule allows it and the zones are of the same zone type (Tap, virtual wire, L2, L3, Tunnel, or External). For example, a security policy rule can allow traffic between two L3 zones, but not between an L3 zone and an L2 zone. The firewall logs traffic that flows between zones when a security policy rule permits interzone traffic.

By default, security policy rules prevent lateral movement of traffic between zones, so malware can't gain access to one zone and then move freely through the network to other targets.



Tunnel zones are for non-encrypted tunnels. You can apply different security policy rules to the tunnel content and to the zone of the outer tunnel, as described in the Tunnel Content Inspection Overview.

© 2018 Palo Alto Networks, Inc.

Zone Defense

Zone protection defends zones from flooding, reconnaissance, packet-based, and protocol-based attacks with zone protection profiles, and from targeted flooding and resource attacks with denial-of-service (DoS) protection profiles and Dos protection policy rules, to complement next-generation firewall features such as App-ID and User-ID. A DoS attack overloads the network with large amounts of unwanted traffic an attempt to disrupt network services.

Unlike security policy rules, there are no default zone protection profiles or DoS protection profiles and DoS protection policy rules. You configure and apply zone protection based on the way you segment your network into zones and on what you want to protect in each zone.

- Zone Defense Tools on page 1067
- How Do the Zone Defense Tools Work? on page 1068
- Zone Protection Profiles on page 1069
- Packet Buffer Protection on page 1071
- DoS Protection Profiles and Policy Rules on page 1071

Zone Defense Tools

Palo Alto Networks firewalls provide three complementary tools to protect the zones in your network:

- **Zone protection profiles** defend the zone at the ingress zone edge against reconnaissance port scan and host sweep attacks, IP packet-based attacks, non-IP protocol attacks, and against flood attacks by limiting the number of connections-per-second of different packet types. The ingress zone is where traffic enters the firewall in the direction of flow from the client to the server (c2s), where the client is the originator of the flow and the server is the responder. The egress zone is where traffic enters the firewall in the direction of flow from the server to the client (s2c).

  Zone protection profiles provide broad defense of the entire zone based on the aggregate traffic entering the zone, protecting against flood attacks and undesirable packet types and options. Zone protection profiles don't control traffic between zones, they control traffic only at the ingress zone. Zone protection profiles don't take individual IP addresses into account because they apply to the aggregate traffic entering the zone (DoS protection policy rules defend individual IP addresses in a zone).

  Use zone protection profiles as a first pass to detect and remove non-compliant traffic. Zone protection profiles defend the network as the session is formed, before the firewall performs DoS protection policy and security policy rule lookups, and consume fewer CPU cycles than a DoS protection policy or security policy rule lookup. If a zone protection profile denies traffic, the firewall doesn't spend CPU cycles on policy rule lookups.
- **DoS protection profiles and DoS protection policy rules** defend against flood attacks and protect specific individual endpoints and resources. The difference between flood protection using a zone protection profile and using a DoS protection profile is that a zone protection profile defends an entire ingress zone based on the aggregate traffic flowing into the zone, while a DoS protection policy rule applies a DoS protection profile that can protect specific IP addresses and address groups, users, zones, and interfaces, so DoS protection is more granular and targeted than a zone protection profile.

  A DoS protection profile sets flood protection thresholds (connections-per-second limits), resource protection thresholds (session limits for specified endpoints and resources), and whether the profile applies to aggregate or classified traffic.

  A DoS protection policy rule specifies:

  - Source, destination, and services match criteria.
  - The action to take when traffic matches the rule.
  - Logging and scheduling options.

© 2018 Palo Alto Networks, Inc.

- The aggregate or classified DoS protection profile the rule applies to matching traffic when protecting resources.

  Aggregate DoS protection profiles and policy rules apply to all of the traffic that matches the specified source, destination, and services. Classified DoS protection profiles and policy rules protect only the traffic that matches the source, destination, or source and destination pair IP addresses and the services specified in the DoS protection policy rule.
- **Security policy rules** affect both the ingress (c2s) and egress (s2c) flows of a session. To establish a session, the incoming traffic must match an existing security policy rule (including the default rules). If there is no match, the firewall discards the packet.

  A Security Policy can protect zones by controlling traffic between zones (interzone) and within zones (intrazone) using criteria including zones, IP addresses, users, applications, services, and URL categories.

  The default security policy rules do not permit traffic to travel between zones, so you need to configure a security rule if you want to allow interzone traffic. All intrazone traffic is allowed by default. You can configure security policy rules to match and control intrazone, interzone, or universal (intrazone and interzone) traffic.

  Zone protection profiles, DoS protection profiles and policy rules, and security policy rules only affect dataplane traffic on the firewall. Traffic originating on the firewall management interface does not cross the dataplane, so the firewall does not match management traffic against these profiles or policy rules.

How Do the Zone Defense Tools Work?

When a packet arrives at the firewall, the firewall attempts to match the packet to an existing session, based on the ingress zone, egress zone, source IP address, destination IP address, protocol, and application derived from the packet header. If the firewall finds a match, then the packet uses the security policy rules that already control the session.

If the packet does not match an existing session, the firewall uses zone protection profiles, DoS protection profiles and policy rules, and security policy rules to determine whether to establish a session or discard the packet, and the level of access the packet receives.

The first protection the firewall applies is the broad edge defense of the zone protection profile, if one exists for the zone. The firewall determines the zone from the interface on which the packet arrives (each interface is assigned to one zone only and all interfaces that carry traffic must belong to a zone). If the zone protection profile denies the packet, the packet is discarded and no DoS protection policy rule or security policy lookup occurs. The firewall applies zone protection profiles only to packets that do not match an existing session. After the firewall establishes a session, the firewall bypasses the zone protection profile lookup for succeeding packets in that session.

The second protection the firewall applies is a DoS protection policy rule lookup. Even if a zone protection profile allows a packet based on the total amount of traffic going to the zone, a DoS protection policy rule and protection profile may deny the packet if it is going to a particular destination or coming from a particular source that has exceeded the flood protection or resource protection settings in the rule's DoS protection profile. If the packet matches a DoS protection policy rule, the firewall applies the rule to the packet. If the rule denies access, the firewall discards the packet and does not perform a security policy lookup. If the rule allows access, the firewall performs a security policy lookup. The DoS protection policy rule is enforced only on new sessions.

The third protection the firewall applies is a Security Policy lookup, which happens only if the zone protection profile and DoS protection policy rules allow the packet. If the firewall finds no security policy rule match for the packet, the firewall discards the packet. If the firewall finds a matching security policy rule, the firewall applies the rule to the packet. The firewall enforces the security policy rule on traffic in both directions (c2s and s2c) for the life of the session.

© 2018 Palo Alto Networks, Inc.

Zone Protection Profiles

Apply a zone protection profile to a zone to defend the entire zone based on the aggregate traffic entering the ingress zone:

- Flood Protection on page 1069
- Reconnaissance Protection on page 1070
- Packet-Based Attack Protection on page 1070
- Protocol Protection on page 1070

Flood Protection

A zone protection profile with flood protection configured defends an entire ingress zone against SYN, ICMP, ICMPv6, UDP, and other IP floods. The firewall measures the aggregate amount of each flood type ingressing the zone in connections-per-second and compares the total to the thresholds configured in the zone protection profile.

For each flood type, you set three thresholds:

- **Alarm Rate**—The number of connections-per-second to trigger an alarm.
- **Activate**—The number of connections-per-second to activate the flood protection mechanism. For ICMP, ICMPv6, UDP, and other IP floods, the protection mechanism is Random Early Drop (RED, also known as Random Early Detection), and packets begin to drop when the number of connections-per-second reaches the **Activate** threshold. For SYN floods, the protection mechanism can be RED or SYN cookies. SYN cookies does not drop packets. As the number of connections-per-second increases above the **Activate** threshold, the firewall drops more packets when RED is the protection mechanism.
- **Maximum**—The number of connections-per-second to drop incoming packets when RED is the protection mechanism.

If the number of connections-per-second exceeds a threshold, the firewall generates an alarm, activates the drop mechanism, or drops all packets when RED is the protection mechanism.

For SYN packets only, you can select **SYN Cookies** instead of dropping the packets with RED. When you use **SYN Cookies**, the firewall acts as a proxy for the target server and responds to the SYN request by generating a SYN-ACK packet and corresponding cookie on behalf of the target. When the firewall receives an ACK packet from the initiator with the correct cookie, the firewall forwards the SYN packet to the target server.

The advantage to using SYN cookies instead of RED is that the firewall drops the offending packets and treats legitimate connections fairly. Because RED randomly drops connections, RED impacts some legitimate traffic. However, using SYN cookies instead of RED uses more firewall resources because the firewall handles the three-way SYN handshake for the target. The tradeoff is using more firewall resources versus not dropping legitimate traffic with RED and offloading the SYN handshake from the target.

Adjust the default threshold values in a zone protection profile to the levels appropriate for your network. The default values are high so that activating a zone protection profile does not unexpectedly drop legitimate traffic.

Adjust the thresholds for your environment by taking a baseline measurement of the peak traffic load for each flood type to determine the normal traffic load for the zone. Set **Alarm Rate** thresholds at 15-20 percent above the baseline number of connections-per-second and monitor the alarms to see if the threshold is reasonable for the legitimate traffic load. Because the normal traffic load experiences some fluctuation, it is best not to drop packets too aggressively.

While determining a baseline and testing the **Alarm Rate** threshold, set the **Activate** and **Maximum** thresholds to a high number to avoid dropping legitimate packets if the thresholds are too aggressive. After you determine a reasonable **Alarm Rate** threshold, set **Activate** and **Maximum** thresholds to drop packets when traffic increases enough beyond normal to indicate a flood attack. Continue to monitor traffic

© 2018 Palo Alto Networks, Inc.

and adjust the thresholds to meet your security objectives and to ensure that the thresholds don't drop legitimate traffic but do prevent unwanted spikes in traffic volume.

A major difference between flood protection using a zone protection profile and a DoS protection profile is where the firewall applies flood protection. Zone protection profiles apply to an entire zone, while DoS protection profiles apply only to the IP addresses, zones, and users specified in the DoS protection policy rule associated with the profile.

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.

Zone protection profiles with reconnaissance protection enabled 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 both 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 *white hat* 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 white hat tests to find and fix network vulnerabilities.

You can set the action to take when reconnaissance traffic (excluding white hat traffic) exceeds the configured threshold when you Configure Reconnaissance Protection on page 1074.

Packet-Based Attack Protection

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

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

You select the drop characteristics for each packet type when you Configure Packet Based Attack Protection.

For example, you can drop malformed IP packets, TCP SYN and SYN-ACK packets that contain data, fragmented ICMP packets, and so on. Each packet type has a set of characteristics and options that you select to control whether the firewall drops a packet. Best Practices for Securing Your Network from Layer4 and Layer 7 Evasions includes some specific recommendations for configuring packet-based attack protection.

Protocol Protection

While packet-based attack protection defends against Layer 3 packet-based attacks, protocol protection defends against non-IP protocol packets. The protocol protection portion of a zone protection profile blocks or allows non-IP protocol packets 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. Configure Protocol Protection to reduce security risks and facilitate regulatory compliance by preventing less secure protocol packets from entering a zone, or an interface in a zone, where they don't belong.

Examples of non-IP protocols that you can block (exclude) or allow (include) are AppleTalk, Banyan VINES, LLDP, NetBEUI, Spanning Tree, and Supervisory Control and Data Acquisition (SCADA) systems such as Generic Object Oriented Substation Event (GOOSE), among many others.

© 2018 Palo Alto Networks, Inc.

You can run App-ID reports to determine whether any non-IP protocol packets are arriving at Layer 2 interfaces on the firewall. Apply the zone protection profile to an ingress security zone for physical interfaces or AE interfaces, thereby controlling interzone traffic (where the protocol packets attempt to enter one zone from another) or intrazone traffic (where the protocol packets traverse a single zone—VLAN—between its interfaces).

Each Include List or Exclude List you configure supports up to 64 Ethertype entries, each identified by its IEEE hexadecimal Ethertype code. Other sources of Ethertype codes are standards.ieee.org/develop/regauth/ethertype/eth.txt and http://www.cavebear.com/archive/cavebear/Ethernet/type.html.

Protocol protection doesn't let you block IPv4 (Ethertype 0x0800), IPv6 (0x86DD), ARP (0x0806), or VLAN-tagged frames (0x8100). These four Ethertypes are always implicitly allowed in an Include List without listing them. They're also implicitly allowed even if you configure an Exclude List; you can't exclude them.

When you configure zone protection for non-IP protocols on zones that have Aggregated Ethernet (AE) interfaces, you can't block or allow a non-IP protocol on only one AE interface because AE interfaces are treated as a group.

Packet Buffer Protection

Packet buffer protection allows you to protect your firewall and network from single session DoS attacks that can overwhelm the firewall's packet buffer and cause legitimate traffic to drop. Although you don't Configure Packet Buffer Protection on page 1079 in a zone protection profile or in a DoS protection profile or policy rule, packet buffer protection defends zones and you enable it when you configure or edit a zone (**Network** > **Zones**).

When you enable packet buffer protection, the firewall monitors sessions from all zones and how each session utilizes the packet buffer. If a session exceeds a configured percentage of packet buffer utilization and traverses an ingress zone with packet buffer protection enabled, then the firewall takes action against that session. The firewall begins by creating an alert log in the System log when a session reaches the first threshold. If a session reaches the second threshold, the firewall mitigates the abuse by implementing Random Early Drop (RED) to throttle the session. If the firewall cannot reduce packet buffer utilization using RED, the Block Hold Time timer begins counting down. When the timer expires, the firewall takes additional mitigation steps (session discard or IP block). The block duration defines how long a session remains discarded or an IP address remains blocked after reaching the block hold time.

In addition to monitoring the buffer utilization of individual sessions, packet buffer protection can also block an IP address if certain criteria are met. While the firewall monitors the packet buffers, if it detects a source IP address rapidly creating sessions that would not individually be seen as an attack, it blocks that IP address.

DoS Protection Profiles and Policy Rules

DoS protection profiles and DoS protection policy rules combine to protect specific areas of your network against packet flood attacks and to protect individual resources against session floods.

DoS protection profiles set the protection thresholds to provide DoS Protection Against Flooding of New Sessions on page 1081 for IP floods (connections-per-second limits), to provide resource protection (maximum concurrent session limits for specified endpoints and resources), and to configure whether the profile applies to aggregate or classified traffic. DoS protection policy rules control where to apply DoS protection and what action to take when traffic matches the criteria defined in the rule.

Unlike a zone protection profile, which protects only the ingress zone, DoS protection profiles and policy rules can protect specific resources inside a zone and traffic flowing between different endpoints and areas. Also unlike a zone protection profile, which supports only aggregate traffic, you can configure aggregate or classified DoS protection profiles and policy rules.

- DoS Protection Policy Rules on page 1072

© 2018 Palo Alto Networks, Inc.

- DoS Protection Profiles on page 1072

DoS Protection Policy Rules

DoS protection policy rules provide granular matching criteria so that you have flexibility in defining what you want to protect:

- Source zone or interface
- Destination zone or interface
- Source IP addresses and address ranges, address group objects, and countries
- Destination IP addresses and address ranges, address group objects, and countries
- Services (by port and protocol)
- Users

The flexible matching criteria enable you to protect entire zones or subnets, a single server, or anything in between. When traffic matches a DoS protection policy rule, the firewall takes one of three actions:

- **Deny**—The firewall denies access and doesn't apply a DoS protection profile. Denying essentially blacklists traffic that matches the rule.
- **Allow**—The firewall permits access and doesn't apply a DoS protection profile. Allowing essentially whitelists traffic that matches the rule.
- **Protect**—The firewall applies the specified DoS protection profile or profiles. A DoS protection policy rule can have one aggregate DoS protection profile and one classified DoS protection profile. Incoming packets count against both DoS protection profiles if the they match the rule. The **Protect** action protects against floods by applying the thresholds set in the DoS protection profile or profiles to traffic that matches the rule.

The firewall only applies DoS protection profiles if the **Action** is **Protect**. If the DoS protection policy rule's **Action** is **Protect**, specify the appropriate aggregate and/or classified DoS protection profile in the rule so that the firewall applies the DoS protection profile to traffic that matches the rule.

You can attach both an aggregate and a classified DoS protection profile to a DoS protection policy rule. The firewall checks and enforces the aggregate rate limits before it checks the classified rate limits, so if the match criteria matches both profiles, the thresholds in the aggregate profile are used first.

DoS Protection Profiles

When you create DoS protection policy rules, you apply DoS protection profiles to the policy rules if the rules have an action of **Protect** (if the action is **Deny** or **Allow**, no DoS protection profile is used).

Configuring flood protection thresholds in a DoS protection profile is similar to configuring Flood Protection on page 1069 in a zone protection profile. The difference is where you apply flood protection. Applying flood protection with a zone protection profile protects the ingress zone, while applying flood protection with a DoS protection profile and policy rule is more granular and targeted, and can even be classified to a single IP address.

For both aggregate and classified DoS protection profiles, as with zone protection profiles, you can:

- Configure SYN, UDP, ICMP, ICMPv6, and other IP flood protection.
- Set alarm, activate, and maximum connections-per-second thresholds. When incoming connections-per-second reach the activate threshold, the firewall begins to drop packets. When the incoming connections-per-second reach the maximum threshold, the firewall drops additional incoming connections.
- Use SYN cookies instead of RED for SYN flood packets.

The advice in zone protection profile Flood Protection on page 1069 about adjusting the default flood threshold values for your network's traffic is valid for setting DoS protection profile flood protection

© 2018 Palo Alto Networks, Inc.

thresholds. Take a baseline measurement of peak traffic loads over a period of time and adjust the flood thresholds to allow the expected legitimate traffic load and to throttle or drop traffic when the load indicates a flood attack. Monitor the traffic and continue to adjust the thresholds until they meet your protection objectives.

Configuring resource protection thresholds in a DoS protection profile sets the maximum number of concurrent sessions that a resource supports. When the number of concurrent sessions reaches its maximum limit, new sessions are dropped. You define the resource you are protecting in a DoS protection policy rule by the resource's source IP address, destination IP address, or the source and destination IP address pair.

An aggregate DoS protection profile applies to all of the traffic that matches the associated DoS protection policy rule, for all sources, destinations, and services allowed for that rule. A classified DoS protection profile can enforce different session rate limits for different groups of end hosts or even for one particular end host.

Here are some examples of what you can do with a classified DoS protection profile:

- To prevent hosts on your network from starting a DoS attack, you can monitor the rate of traffic each host in a source address group initiates. To do this, set an appropriate alarm threshold in a DoS protection profile to notify you if a host initiates an unusually large amount of traffic, and create a DoS protection policy rule that applies the profile to the source address group. Investigate any hosts that initiate enough traffic to set off the alarm.
- To protect critical web or DNS servers on your network, protect the individual servers. To do this, set appropriate flooding and resource protection thresholds in a DoS protection profile, and create a DoS protection policy rule that applies the profile to each server's IP address by adding the IP addresses as the rule's destination criteria.
- Track the flow between a pair of endpoints by setting appropriate thresholds in the DoS protection profile and creating a DoS protection policy rule that specifies the source and destination IP addresses of the endpoints as the matching criteria.



Do not use source IP classification for internet-facing zones in classified DoS protection policy rules. The firewall does not have the capacity to store counters for every possible IP address on the internet.

© 2018 Palo Alto Networks, Inc.

Configure Zone Protection to Increase Network Security

The following topics provide zone protection configuration examples:

- Configure Reconnaissance Protection on page 1074
- Configure Packet Based Attack Protection on page 1074
- Configure Protocol Protection on page 1075
- Configure Packet Buffer Protection on page 1079

Configure Reconnaissance Protection

Configure one of the following Reconnaissance Protection on page 1070 actions for the firewall to take in response to the corresponding reconnaissance attempt:

- **Allow**—The firewall allows the port scan or host sweep reconnaissance to continue.
- **Alert**—The firewall generates an alert for each port scan or host sweep that matches the configured threshold within the specified time interval. Alert is the default action.
- **Block**—The firewall drops all subsequent packets from the source to the destination for the remainder of the specified time interval.
- **Block IP**—The firewall drops all subsequent packets for the specified **Duration**, in seconds (the range is 1-3,600). **Track By** determines whether the firewall blocks source or source-and-destination traffic.

STEP 1 | Configure Reconnaissance Protection.

1. Select **Network** > **Network Profiles** > **Zone Protection**.
2. Select a Zone Protection profile or **Add** a new profile and enter a **Name** for it.
3. On the Reconnaissance Protection tab, select the scan types to protect against.
4. Select an **Action** for each scan. If you select Block IP, you must also configure **Track By** (source or source-and-destination) and **Duration**.
5. Set the **Interval** in seconds. This options defines the time interval for port scan and host sweep detection.
6. Set the **Threshold**. The threshold defines the number of port scan events or host sweeps that occurs within the interval configured above that triggers an action.

STEP 2 | (Optional) Configure a Source Address Exclusion.

1. On the Reconnaissance Protection tab, **Add** a Source Address Exclusion.
    1. Enter a descriptive **Name** for the whitelisted address.
    2. Set the Address Type to **IPv4** or **IPv6** and then select an address object or enter an IP address.
    3. Click **OK**.
2. Click **OK** to save the Zone Protection profile.
3. **Commit** your changes.

Configure Packet Based Attack Protection

To enhance security for a zone, Packet-Based Attack Protection on page 1070 allows you to specify whether the firewall drops IP, IPv6, TCP, ICMP, or ICMPv6 packets that have certain characteristics or strips certain options from the packets.

© 2018 Palo Alto Networks, Inc.

For example, you can drop TCP SYN and SYN-ACK packets that contain data in the payload during a TCP three-way handshake. A Zone Protection profile by default is set to drop SYN and SYN-ACK packets with data (you must apply the profile to the zone).

The TCP Fast Open option (RFC 7413) preserves the speed of a connection setup by including data in the payload of SYN and SYN-ACK packets. A Zone Protection profile treats handshakes that use the TCP Fast Open option separately from other SYN and SYN-ACK packets; the profile by default is set to allow the handshake packets if they contain a valid Fast Open cookie.



If you have existing Zone Protection profiles in place when you upgrade to PAN-OS 8.0, the three default settings will apply to each profile and the firewall will act accordingly.

STEP 1 | Create a Zone Protection profile for packet based attack protection.

1. Select **Network** > **Network Profiles** > **Zone Protection** and **Add** a new profile.
2. Enter a **Name** for the profile and an optional **Description**.
3. Select **Packet Based Attack Protection**.
4. On each tab (**IP Drop**, **TCP Drop**, **ICMP Drop**, **IPv6 Drop**, and **ICMPv6 Drop**), select the settings you want to enforce to protect a zone.
5. Click **OK**.

STEP 2 | Apply the Zone Protection profile to a security zone that is assigned to interfaces you want to protect.

1. Select **Network** > **Zones** and select the zone where you want to assign the Zone Protection profile.
2. **Add** the **Interfaces** belonging to the zone.
3. For **Zone Protection Profile**, select the profile you just created.
4. Click **OK**.

STEP 3 | Commit.

Click **Commit**.

Configure Protocol Protection

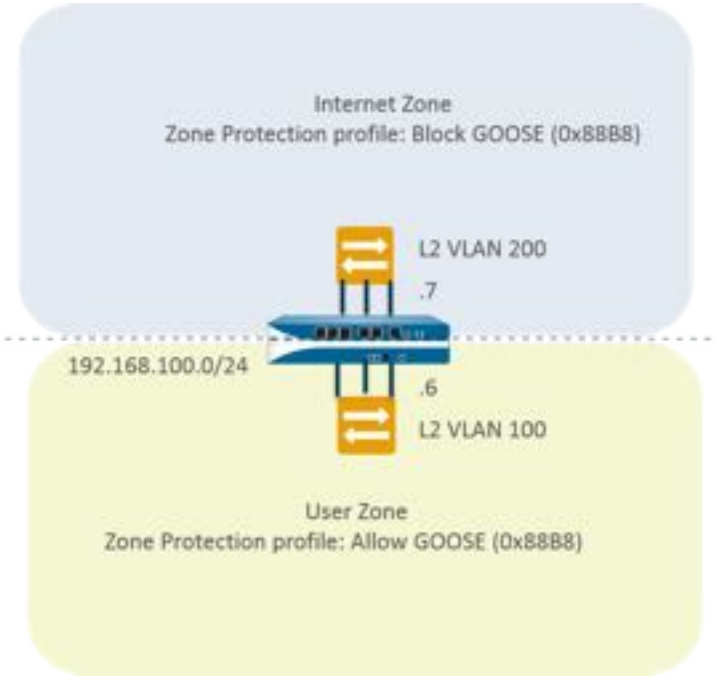
Protect virtual wire or Layer 2 security zones from non-IP protocol packets by using Protocol Protection on page 1070.

- Use Case: Non-IP Protocol Protection Between Security Zones on Layer 2 Interfaces on page 1075
- Use Case: Non-IP Protocol Protection Within a Security Zone on Layer 2 Interfaces on page 1077

Use Case: Non-IP Protocol Protection Between Security Zones on Layer 2 Interfaces

In this use case, the firewall is in a Layer 2 VLAN divided into two subinterfaces. VLAN 100 is 192.168.100.1/24, subinterface .6. VLAN 200 is 192.168.100.1/24, subinterface .7. Non-IP protocol protection applies to ingress zones. In this use case, if the Internet zone is the ingress zone, the firewall blocks the Generic Object Oriented Substation Event (GOOSE) protocol. If the User zone is the ingress zone, the firewall allows the GOOSE protocol. The firewall implicitly allows IPv4, IPv6, ARP, and VLAN-tagged frames in both zones.

© 2018 Palo Alto Networks, Inc.



STEP 1 | Configure two VLAN subinterfaces.

1. Select **Network** > **Interfaces** > **VLAN** and **Add** an interface.
2. **Interface Name** defaults to vlan. After the period, enter 7.
3. On the **Config** tab, **Assign Interface To** the **VLAN** 200.
4. Click **OK**.
5. Select **Network** > **Interfaces** > **VLAN** and **Add** an interface.
6. **Interface Name** defaults to vlan. After the period, enter 6.
7. On the **Config** tab, **Assign Interface To** the **VLAN** 100.
8. Click **OK**.

STEP 2 | Configure protocol protection in a Zone Protection profile to block GOOSE protocol packets.

1. Select **Network** > **Network Profiles** > **Zone Protection** and **Add** a profile.
2. Enter the **Name** Block GOOSE.
3. Select **Protocol Protection**.
4. Choose **Rule Type** of **Exclude List**.
5. Enter the **Protocol Name**, GOOSE, to easily identify the Ethertype on the list. The firewall doesn't verify that the name you enter matches the Ethertype code; it uses only the Ethertype code to filter.
6. Enter **Ethertype** code 0x88B8. The Ethertype must be preceded by 0x to indicate a hexadecimal value. Range is 0x0000 to 0xFFFF.
7. Select **Enable** to enforce the protocol protection. You can disable a protocol on the list, for example, for testing.
8. Click **OK**.

STEP 3 | Apply the Zone Protection profile to the Internet zone.

1. Select **Network** > **Zones** and **Add** a zone.
2. Enter the **Name** of the zone, Internet.
3. For **Location**, select the virtual system where the zone applies.
4. For **Type**, select **Layer2**.
5. **Add** the **Interface** that belongs to the zone, vlan.7.
6. For **Zone Protection Profile**, select the profile Block GOOSE.

© 2018 Palo Alto Networks, Inc.

7. Click **OK**.

STEP 4 | Configure protocol protection to allow GOOSE protocol packets.

Create another Zone protection profile named Allow GOOSE, and choose **Rule Type** of **Include List**.



When configuring an Include list, include all required non-IP protocols; an incomplete list can result in legitimate non-IP traffic being blocked.

STEP 5 | Apply the Zone Protection profile to the User zone.

1. Select **Network** > **Zones** and **Add** a zone.
2. Enter the **Name** of the zone, User.
3. For **Location**, select the virtual system where the zone applies.
4. For **Type**, select **Layer2**.
5. **Add** the **Interface** that belongs to the zone, vlan.6.
6. For **Zone Protection Profile**, select the profile Allow GOOSE.
7. Click **OK**.

STEP 6 | Commit.

Click **Commit**.

STEP 7 | View the number of non-IP packets the firewall has dropped based on protocol protection.

Access the CLI.

```
> show counter global name pkt_nonip_pkt_drop
> show counter global name pkt_nonip_pkt_drop delta yes
```

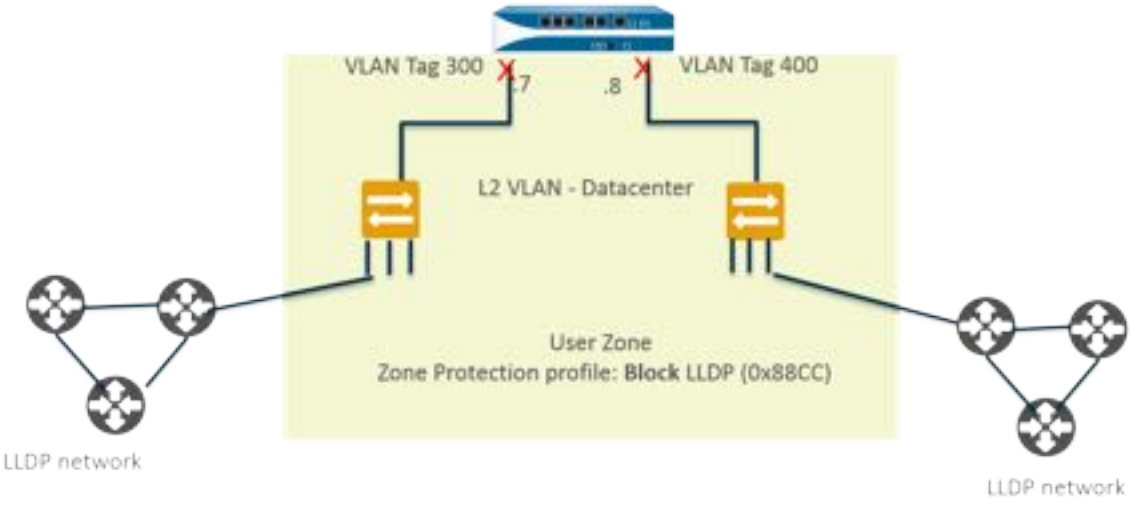
Use Case: Non-IP Protocol Protection Within a Security Zone on Layer 2 Interfaces

If you don't implement a Zone Protection profile with non-IP protocol protection, the firewall allows non-IP protocols in a single zone to go from one Layer 2 interface to another. In this use case, blacklisting LLDP packets ensures that LLDP for one network doesn't discover a network reachable through another interface in the zone.

In the following figure, the Layer 2 VLAN named Datacenter is divided into two subinterfaces: 192.168.1.1/24, subinterface .7 and 192.168.1.2/24, subinterface .8. The VLAN belongs to the User zone. By applying a Zone Protection profile that blocks LLDP to the User zone:

- Subinterface .7 blocks LLDP from its switch to the firewall at the red X on the left, preventing that traffic from reaching subinterface .8.
- Subinterface .8 blocks LLDP from its switch to the firewall at the red X on the right, preventing that traffic from reaching subinterface .7.

© 2018 Palo Alto Networks, Inc.



STEP 1 | Create a subinterface for an Ethernet interface.

1. Select **Network** > **Interfaces** > **Ethernet** and select a Layer 2 interface, in this example, ethernet1/1.
2. Select **Add Subinterfaces**.
3. The **Interface Name** defaults to the interface (ethernet 1/1). After the period, enter 7.
4. For **Tag**, enter 300.
5. For **Security Zone**, select User.
6. Click **OK**.

STEP 2 | Create a second subinterface for the Ethernet interface.

1. Select **Network** > **Interfaces** > **Ethernet** and select the Layer 2 interface: ethernet1/1.
2. Select **Add Subinterfaces**.
3. The **Interface Name** defaults to the interface (ethernet 1/1). After the period, enter 8.
4. For **Tag**, enter 400.
5. For **Security Zone**, select User.
6. Click **OK**.

STEP 3 | Create a VLAN for the Layer2 interface and two subinterfaces.

1. Select **Network** > **VLANs** and **Add** a VLAN.
2. Enter the **Name** of the VLAN; for this example, enter Datacenter.
3. For **VLAN Interface**, select **None**.
4. For **Interfaces**, click **Add** and select the Layer 2 interface: ethernet1/1, and two subinterfaces: ethernet1/1.7 and ethernet1/1.8.
5. Click **OK**.

STEP 4 | Block non-IP protocol packets in a Zone Protection profile.

1. Select **Network** > **Network Profiles** > **Zone Protection** and **Add** a profile.
2. Enter the **Name**, in this example, Block LLDP.
3. Enter a profile **Description**—Block LLDP packets from an LLDP network to other interfaces in the zone (intrazone).
4. Select **Protocol Protection**.
5. Choose **Rule Type** of **Exclude List**.
6. Enter **Protocol Name** LLDP.
7. Enter **Ethertype** code 0x88cc. The Ethertype must be preceded by 0x to indicate a hexadecimal value.

© 2018 Palo Alto Networks, Inc.

8. Select **Enable**.
9. Click **OK**.

STEP 5 | Apply the Zone Protection profile to the security zone to which Layer 2 VLAN belongs.

1. Select **Network** > **Zones**.
2. **Add** a zone.
3. Enter the **Name** of the zone, User.
4. For **Location**, select the virtual system where the zone applies.
5. For **Type**, select **Layer2**.
6. **Add** an **Interface** that belongs to the zone, ethernet1/1.7
7. **Add** an **Interface** that belongs to the zone, ethernet1/1.8.
8. For **Zone Protection Profile**, select the profile Block LLDP.
9. Click **OK**.

STEP 6 | Commit.

Click **Commit**.

STEP 7 | View the number of non-IP packets the firewall has dropped based on protocol protection.

Access the CLI.

```
> show counter global name pkt_nonip_pkt_drop
> show counter global name pkt_nonip_pkt_drop delta yes
```

Configure Packet Buffer Protection

You configure Packet Buffer Protection settings globally and then apply them per ingress zone. When the firewall detects high buffer utilization, the firewall only monitors and takes action against sessions from zones with packet buffer protection enabled. Therefore, if the abusive session is from a zone without packet buffer protection, the high packet buffer utilization continues. Packet buffer protection can be applied to a zone but it is not active until global settings are configured and enabled.

STEP 1 | Configure the global session thresholds.

1. Select **Device** > **Setup** > **Session**.
2. **Edit** the Session Settings.
3. Select the **Packet Buffer Protection** check box to enable and configure the packet buffer protection thresholds.
4. Enter a value for each threshold and timer to define the packet buffer protection behavior.
   - **Alert (%)**—When packet buffer utilization exceeds this threshold for more than 10 seconds, the firewall creates a log event every minute. The firewall generates log events when packet buffer protection is enabled globally. The default threshold is 50% and the range is 0% to 99%. If the value is 0%, the firewall does not create a log event.
   - **Activate (%)**—When a packet buffer utilization exceeds this threshold, the firewall applies RED to abusive sessions. The default threshold is 50% and the range is 0% to 99%. If the value is 0%, the firewall does not apply RED.

     The firewall records alert events in the System log and events for dropped traffic, discarded sessions, and blocked IP address in the Threat log.
   - **Block Hold Time (sec)**—The amount of time a RED-mitigated session is allowed to continue before the firewall discards it. By default, the block hold time is 60 seconds. The range is 0 to 65,535 seconds. If the value is 0, the firewall does not discard sessions based on packet buffer protection.

© 2018 Palo Alto Networks, Inc.

- **Block Duration (sec)**—This setting defines how long a session remains discarded or an IP address remains blocked. The default is 3,600 seconds with a range of 1 seconds to 15,999,999 seconds.

5. Click **OK**.
6. **Commit** your changes.

STEP 2 | Enable packet buffer protection on an ingress zone.

1. Select **Network** > **Zones**.
2. Choose an ingress zone and click on its name.
3. Select the **Enable Packet Buffer Protection** check box in the Zone Protection section.
4. Click **OK**.
5. **Commit** your changes.

© 2018 Palo Alto Networks, Inc.

DoS Protection Against Flooding of New Sessions

DoS protection against flooding of new sessions is beneficial against high-volume single-session and multiple-session attacks. In a single-session attack, an attacker uses a single session to target a device behind the firewall. If a Security rule allows the traffic, the session is established and the attacker initiates an attack by sending packets at a very high rate with the same source IP address and port number, destination IP address and port number, and protocol, trying to overwhelm the target. In a multiple-session attack, an attacker uses multiple sessions (or connections per second [cps]) from a single host to launch a DoS attack.

This feature defends against DoS attacks of new sessions only, that is, traffic that has not been offloaded to hardware. An offloaded attack is not protected by this feature. However, this topic describes how you can create a Security policy rule to reset the client; the attacker reinitiates the attack with numerous connections per second and is blocked by the defenses illustrated in this topic.

DoS Protection Profiles and Policy Rules on page 1071 work together to provide protection against flooding of many incoming SYN, UDP, ICMP, and ICMPv6 packets, and other types of IP packets. You determine what thresholds constitute flooding. In general, the DoS Protection profile sets the thresholds at which the firewall generates a DoS alarm, takes action such as Random Early Drop, and drops additional incoming connections. A DoS Protection policy rule that is set to protect (rather than to allow or deny packets) determines the criteria for packets to match (such as source address) in order to be counted toward the thresholds. This flexibility allows you to blacklist certain traffic, or whitelist certain traffic and treat other traffic as DoS traffic. When the incoming rate exceeds your maximum threshold, the firewall blocks incoming traffic from the source address.

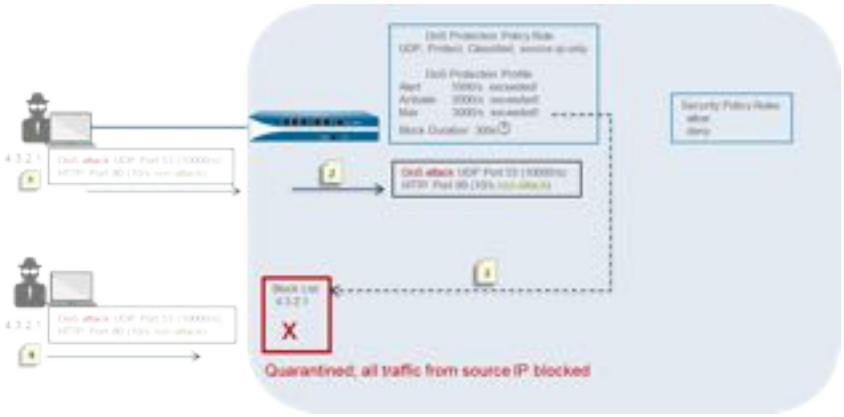
- Multiple-Session DoS Attack on page 1081
- Single-Session DoS Attack on page 1084
- Configure DoS Protection Against Flooding of New Sessions on page 1084
- End a Single Session DoS Attack on page 1086
- Identify Sessions That Use an Excessive Percentage of the Packet Buffer on page 1087
- Discard a Session Without a Commit on page 1089

Multiple-Session DoS Attack

Configure DoS Protection Against Flooding of New Sessions by configuring a DoS Protection policy rule, which determines the criteria that, when matched by incoming packets, trigger the **Protect** action. The DoS Protection profile counts each new connection toward the Alarm Rate, Activate Rate, and Max Rate thresholds. When the incoming new connections per second exceed the Activate Rate, the firewall takes the action specified in the DoS Protection profile.

The following figure and table describe how the Security policy rules, DoS Protection policy rules and profile work together in an example.

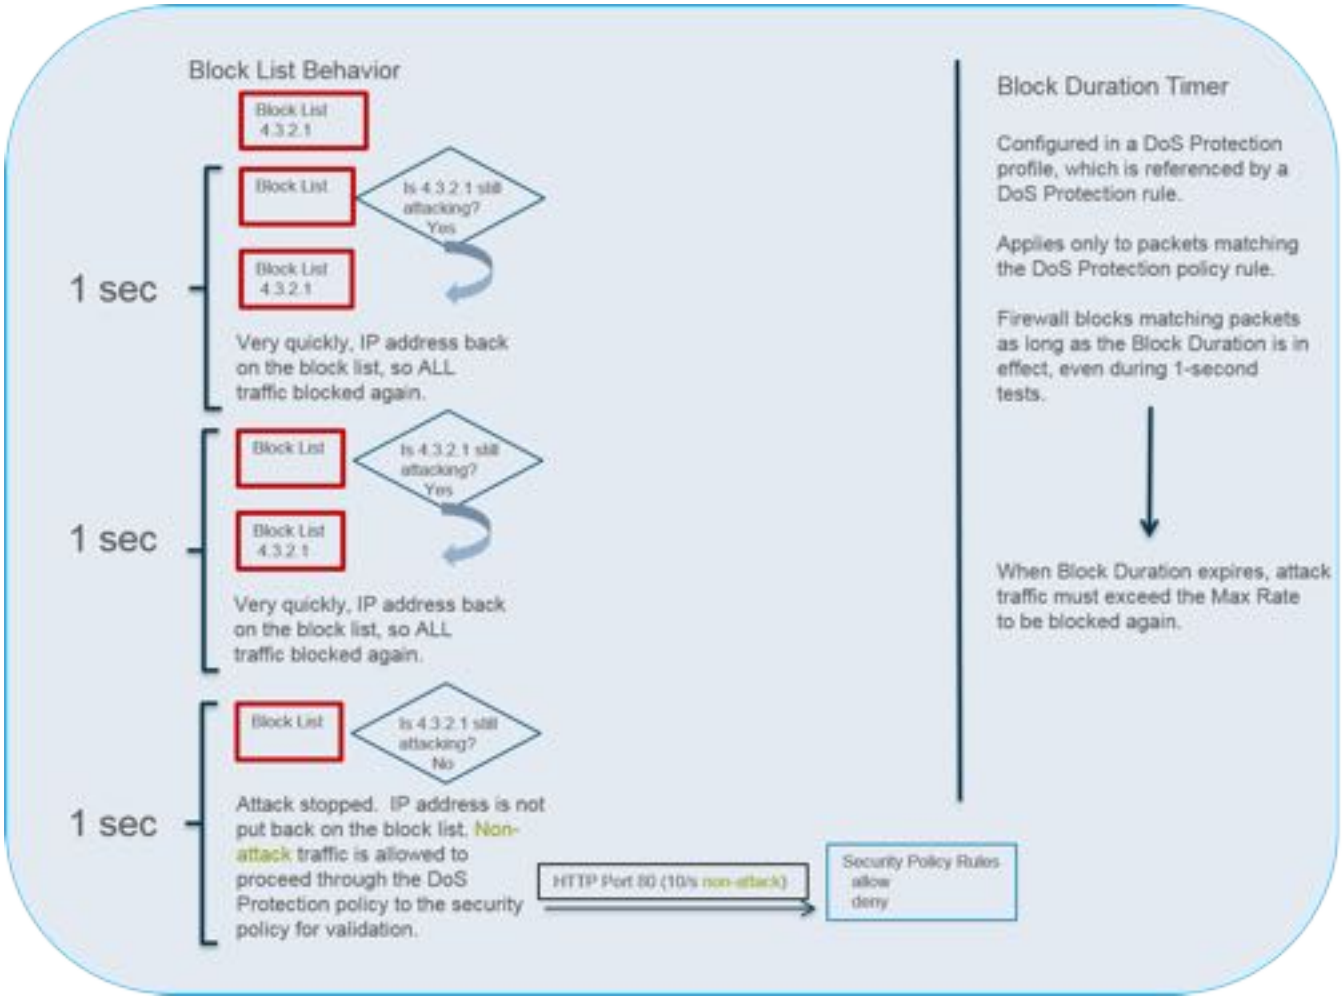
© 2018 Palo Alto Networks, Inc.



Sequence of Events as Firewall Quarantines an IP Address	
1	In this example, an attacker launches a DoS attack at a rate of 10,000 new connections per second to UDP port 53. The attacker also sends 10 new connections per second to HTTP port 80.
2	The new connections match criteria in the DoS Protection policy rule, such as a source zone or interface, source IP address, destination zone or interface, destination IP address, or a service, among other settings. In this example, the policy rule specifies UDP. The DoS Protection policy rule also specifies the **Protect** action and **Classified**, two settings that dynamically put the DoS Protection profile settings into effect. The DoS Protection profile specifies that a Max Rate of 3000 packets per second is allowed. When incoming packets match the DoS Protection policy rule, new connections per second are counted toward the **Alert**, **Activate**, and **Max Rate** thresholds. *You can also use a Security policy rule to block all traffic from the source IP address if you deem that address to be malicious all the time.*
3	The 10,000 new connections per second exceed the **Max Rate** threshold. When all of the following occur: • the threshold is exceeded, • a **Block Duration** is specified, and • **Classified** is set to include source IP address, the firewall puts the offending source IP address on the block list.
4	An IP address on the block list is in quarantine, meaning all traffic from that IP address is blocked. The firewall blocks the offending source IP address before additional attack packets reach the Security policy.

The following figure describes in more detail what happens after an IP address that matches the DoS Protection policy rule is put on the block list. It also describes the Block Duration timer.

© 2018 Palo Alto Networks, Inc.



Every one second, the firewall allows the IP address to come off the block list so that the firewall can test the traffic patterns and determine if the attack is ongoing. The firewall takes the following action:

- During this one-second test period, the firewall allows packets that don't match the DoS Protection policy criteria (HTTP traffic in this example) through the DoS Protection policy rules to the Security policy for validation. Very few packets, if any, have time to get through because the first attack packet that the firewall receives after the IP address is let off the block list will match the DoS Protection policy criteria, quickly causing the IP address to be placed back on the block list for another second. The firewall repeats this test each second until the attack stops.
- The firewall blocks all attack traffic from going past the DoS Protection policy rules (the address remains on the block list) until the Block Duration expires.

When the attack stops, the firewall does not put the IP address back on the block list. The firewall allows non-attack traffic to proceed through the DoS Protection policy rules to the Security policy rules for evaluation. You must configure a Security policy rule to allow or deny traffic because without one, an implicit Deny rule denies all traffic.

The block list is based on a source zone and source address combination. This behavior allows duplicate IP addresses to exist as long as they are in different zones belonging to separate virtual routers.

The Block Duration setting in a DoS Protection profile specifies how long the firewall blocks the [offending] packets that match a DoS Protection policy rule. The attack traffic remains blocked until the Block Duration expires, after which the attack traffic must again exceed the Max Rate threshold to be blocked again.

© 2018 Palo Alto Networks, Inc.

If the attacker uses multiple sessions or bots that initiate multiple attack sessions, the sessions count toward the thresholds in the DoS Protection profile without a Security policy deny or drop rule in place. Hence, a single-session attack requires a Security policy deny or drop rule in order for each packet to count toward the thresholds; a multiple-session attack does not.

Therefore, the DoS protection against flooding of new sessions allows the firewall to efficiently defend against a source IP address while attack traffic is ongoing and to permit non-attack traffic to pass as soon as the attack stops. Putting the offending IP address on the block list allows the DoS protection functionality to take advantage of the block list, which is designed to quarantine all activity from that source IP address, such as packets with a different application. Quarantining the IP address from all activity protects against a modern attacker who attempts a rotating application attack, in which the attacker simply changes applications to start a new attack or uses a combination of different attacks in a hybrid DoS attack. You can Monitor Blocked IP Addresses to view the block list, remove entries from it, and get additional information about an IP address on the block list.

Beginning with PAN-OS 7.0.2, it is a change in behavior that the firewall places the attacking source IP address on the block list. When the attack stops, non-attack traffic is allowed to proceed to Security policy enforcement. The attack traffic that matched the DoS Protection profile and DoS Protection policy rules remains blocked until the Block Duration expires.

Single-Session DoS Attack

A single-session DoS attack typically will not trigger Zone or DoS Protection profiles because they are attacks that are formed after the session is created. These attacks are allowed by the Security policy because a session is allowed to be created, and after the session is created, the attack drives up the packet volume and takes down the target device.

Configure DoS Protection Against Flooding of New Sessions on page 1084 to protect against flooding of new sessions (single-session and multiple-session flooding). In the event of a single-session attack that is underway, additionally End a Single Session DoS Attack on page 1086.

Configure DoS Protection Against Flooding of New Sessions

STEP 1 | Configure Security policy rules to deny traffic from the attacker's IP address and allow other traffic based on your network needs. You can specify any of the match criteria in a Security policy rule, such as source IP address. (Required for single-session attack mitigation or attacks that have not triggered the DoS Protection policy threshold; optional for multiple-session attack mitigation).

This step is one of the steps typically performed to stop an existing attack. See End a Single Session DoS Attack on page 1086.

- Create a Security Policy Rule on page 991

STEP 2 | Configure a DoS Protection profile for flood protection.

Because flood attacks can occur over multiple protocols, as a best practice, activate protection for all of the flood types in the DoS Protection profile.

1. Select **Objects** > **Security Profiles** > **DoS Protection** and **Add** a profile **Name**.
2. Select **Classified** as the **Type**.
3. For **Flood Protection**, select all types of flood protection:
   - **SYN Flood**

© 2018 Palo Alto Networks, Inc.

- **UDP Flood**
- **ICMP Flood**
- **ICMPv6 Flood**
- **Other IP Flood**

4. When you enable **SYN Flood**, select the **Action** that occurs when connections per second (cps) exceed the **Activate Rate** threshold:
    1. **Random Early Drop**—The firewall uses an algorithm to progressively start dropping that type of packet. If the attack continues, the higher the incoming cps rate (above the **Activate Rate**) gets, the more packets the firewall drops. The firewall drops packets until the incoming cps rate reaches the **Max Rate**, at which point the firewall drops all incoming connections. **Random Early Drop** (RED) is the default action for **SYN Flood**, and the only action for **UDP Flood**, **ICMP Flood**, **ICMPv6 Flood**, and **Other IP Flood**. RED is more efficient than SYN Cookies and can handles larger attacks, but doesn't discern between good and bad traffic.
    2. **SYN Cookies**—Rather than immediately sending the SYN to the server, the firewall generates a cookie (on behalf of the server) to send in the SYN-ACK to the client. The client responds with its ACK and the cookie; upon this validation the firewall then sends the SYN to the server. The **SYN Cookies** action requires more firewall resources than **Random Early Drop**; it's more discerning because it affects bad traffic.
5. (Optional) On each of the flood tabs, change the following thresholds to suit your environment:
    - **Alarm Rate (connections/s)**—Specify the threshold rate (cps) above which a DoS alarm is generated. (Range is 0-2,000,000; default is 10,000.)
    - **Activate Rate (connections/s)**—Specify the threshold rate (cps) above which a DoS response is activated. When the **Activate Rate** threshold is reached, **Random Early Drop** occurs. Range is 0-2,000,000; default is 10,000. (For SYN Flood, you can select the action that occurs.)
    - **Max Rate (connections/s)**—Specify the threshold rate of incoming connections per second that the firewall allows. When the threshold is exceeded, new connections that arrive are dropped. (Range is 2-2,000,000; default is 40,000.)

    The default threshold values in this step are only starting points and might not be appropriate for your network. You must analyze the behavior of your network to properly set initial threshold values.
6. On each of the flood tabs, specify the **Block Duration** (in seconds), which is the length of time the firewall blocks packets that match the DoS Protection policy rule that references this profile. Specify a value greater than zero. (Range is 1-21,600; default is 300.)

    Set a low Block Duration value if you are concerned that packets you incorrectly identify as attack traffic will be blocked unnecessarily.

    Set a high **Block Duration** value if you are more concerned about blocking volumetric attacks than you are about incorrectly blocking packets that aren't part of an attack.
7. Click **OK**.

STEP 3 | Configure a DoS Protection policy rule that specifies the criteria for matching the incoming traffic.



The firewall resources are finite, so you wouldn't want to classify using source address on an internet-facing zone because there can be an enormous number of unique IP addresses that match the DoS Protection policy rule. That would require many counters and the firewall would run out of tracking resources. Instead, define a DoS Protection policy rule that classifies using the destination address (of the server you are protecting).

© 2018 Palo Alto Networks, Inc.

1. Select **Policies** > **DoS Protection** and **Add** a **Name** on the **General** tab. The name is case-sensitive and can be a maximum of 31 characters, including letters, numbers, spaces, hyphens, and underscores.
2. On the **Source** tab, choose the **Type** to be a **Zone** or **Interface**, and then **Add** the zone(s) or interface(s). Choose zone or interface depending on your deployment and what you want to protect. For example, if you have only one interface coming into the firewall, choose Interface.
3. (Optional) For **Source Address**, select **Any** for any incoming IP address to match the rule or **Add** an address object such as a geographical region.
4. (Optional) For **Source User**, select **any** or specify a user.
5. (Optional) Select **Negate** to match any sources except those you specify.
6. (Optional) On the **Destination** tab, choose the **Type** to be a **Zone** or **Interface**, and then **Add** the destination zone(s) or interface(s). For example, enter the security zone you want to protect.
7. (Optional) For **Destination Address**, select **Any** or enter the IP address of the device you want to protect.
8. (Optional) On the **Option/Protection** tab, **Add** a **Service**. Select a service or click **Service** and enter a **Name**. Select **TCP** or **UDP**. Enter a **Destination Port**. Not specifying a particular service allows the rule to match a flood of any protocol type without regard to an application-specific port.
9. On the **Option/Protection** tab, for **Action**, select **Protect**.
10. Select **Classified**.
11. For **Profile**, select the name of the **DoS Protection** profile you created.
12. For **Address**, select **source-ip-only** or **src-dest-ip-both**, which determines the type of IP address to which the rule applies. Choose the setting based on how you want the firewall to identify offending traffic:
    - Specify **source-ip-only** if you want the firewall to classify only on the source IP address. Because attackers often test the entire network for hosts to attack, **source-ip-only** is the typical setting for a wider examination.
    - Specify **src-dest-ip-both** if you want to protect against DoS attacks only on the server that has a specific destination address, and you also want to ensure that every source IP address won't surpass a specific cps threshold to that server.
13. Click **OK**.

STEP 4 | Commit.

Click **Commit**.

End a Single Session DoS Attack

To mitigate a single-session DoS attack, you would still Configure DoS Protection Against Flooding of New Sessions in advance. At some point after you configure the feature, a session might be established before you realize a DoS attack (from the IP address of that session) is underway. When you see a single-session DoS attack, perform the following task to end the session, so that subsequent connection attempts from that IP address trigger the DoS protection against flooding of new sessions.

STEP 1 | Identify the source IP address that is causing the attack.

For example, use the firewall Packet Capture feature with a destination filter to collect a sample of the traffic going to the destination IP address. Alternatively, use the ACC to filter on destination address to view the activity to the target host being attacked.

STEP 2 | Create a DoS Protection policy rule that will block the attacker's IP address after the attack thresholds are exceeded.

STEP 3 | Create a Security policy rule to deny the source IP address and its attack traffic.

© 2018 Palo Alto Networks, Inc.

STEP 4 | End any existing attacks from the attacking source IP address by executing the **clear session all filter source *<ip-address>*** operational command.

Alternatively, if you know the session ID, you can execute the **clear session id *<value>*** command to end that session only.

*If you use the **clear session all filter source *<ip-address>*** command, all sessions matching the source IP address are discarded, which can include both good and bad sessions.*

After you end the existing attack session, any subsequent attempts to form an attack session are blocked by the Security policy. The DoS Protection policy counts all connection attempts toward the thresholds. When the Max Rate threshold is exceeded, the source IP address is blocked for the Block Duration, as described in Multiple-Session DoS Attack.

Identify Sessions That Use an Excessive Percentage of the Packet Buffer

When a firewall exhibits signs of resource depletion, it might be experiencing an attack that is sending an overwhelming number of packets. In such events, the firewall starts buffering inbound packets. You can quickly identify the sessions that are using an excessive percentage of the packet buffer and mitigate their impact by discarding them.

Perform the following task on any hardware-based firewall model (not a VM-Series firewall) to identify, for each slot and dataplane, the packet buffer percentage used, the top five sessions using more than two percent of the packet buffer, and the source IP addresses associated with those sessions. Having that information allows you to take appropriate action.

STEP 1 | View firewall resource usage, top sessions, and session details. Execute the following operational command in the CLI (sample output from the command follows):

```
admin@PA-7050> show running resource-monitor ingress-backlogs
-- SLOT:s1, DP:dp1 -- USAGE - ATOMIC: 92%  TOTAL: 93%
TOP SESSIONS:SESS-ID      PCT   GRP-ID   COUNT
6            92%   1        156                  7        1732
SESSION DETAILS SESS-ID PROTO SZONESRC       SPORT  DST       DPORT  IGR-
IF    EGR-IF       APP
6    6     trust 192.168.2.35 55653  10.1.8.89 80  ethernet1/21 ethernet1/22
 undecided
```

The command displays a maximum of the top five sessions that each use 2% or more of the packet buffer.

The sample output above indicates that Session 6 is using 92% of the packet buffer with TCP packets (protocol 6) coming from source IP address 192.168.2.35.

- **SESS-ID**—Indicates the global session ID that is used in all other **show session** commands. The global session ID is unique within the firewall.
- **GRP-ID**—Indicates an internal stage of processing packets.
- **COUNT**—Indicates how many packets are in that GRP-ID for that session.
- **APP**—Indicates the App-ID extracted from the Session information, which can help you determine whether the traffic is legitimate. For example, if packets use a common TCP or UDP port but the CLI output indicates an APP of `undecided`, the packets are possibly attack traffic. The APP is `undecided` when Application IP Decoders cannot get enough information to determine the application. An APP of `unknown` indicates that Application IP Decoders cannot determine the

© 2018 Palo Alto Networks, Inc.

application; a session of `unknown` APP that uses a high percentage of the packet buffer is also suspicious.

To restrict the display output:

On a PA-7000 Series model only, you can limit output to a slot, a dataplane, or both. For example:

```
admin@PA-7050> show running resource-monitor ingress-backlogs slot s1
admin@PA-7050> show running resource-monitor ingress-backlogs slot s1 dp
 dp1
```

On PA-5000 Series, PA-5200 Series, and PA-7000 Series models only, you can limit output to a dataplane. For example:

```
admin@PA-5060> show running resource-monitor ingress-backlogs dp dp1
```

STEP 2 | Use the command output to determine whether the source at the source IP address using a high percentage of the packet buffer is sending legitimate or attack traffic.

In the sample output above, a single-session attack is likely occurring. A single session (Session ID 6) is using 92% of the packet buffer for Slot 1, DP 1, and the application at that point is `undecided`.

- If you determine a single user is sending an attack and the traffic is not offloaded, you can End a Single Session DoS Attack. At a minimum, you can Configure DoS Protection Against Flooding of New Sessions.
- On a hardware model that has a field-programmable gate array (FPGA), the firewall offloads traffic to the FPGA when possible to increase performance. If the traffic is offloaded to hardware, clearing the session does not help because then it is the software that must handle the barrage of packets. You should instead Discard a Session Without a Commit.

To see whether a session is offloaded or not, use the **`show session id <session-id>`** operational command in the CLI as shown in the following example. The `layer7 processing` value indicates `completed` for sessions offloaded or `enabled` for sessions not offloaded.

© 2018 Palo Alto Networks, Inc.

```
admin@PA-5060> show session id 68088184

Session        68088184

        c2s flow:
                source:      1.1.42.15 [trust]
                dst:         1.2.27.99
                proto:       6
                sport:       55993           dport:      6881
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
                offload:     Yes

        s2c flow:
                source:      1.2.27.99 [untrust]
                dst:         1.1.42.15
                proto:       6
                sport:       6881            dport:      55993
                state:       ACTIVE          type:       FLOW
                src user:    unknown
                dst user:    unknown
                offload:     Yes

        DP                                   : 2
        index(local):                        : 979320
        start time                           : Tue Oct 27 14:20:09 2015
        timeout                              : 1200 sec
        time to live                         : 1167 sec
        total byte count(c2s)                : 270
        total byte count(s2c)                : 270
        layer7 packet count(c2s)             : 3
        layer7 packet count(s2c)             : 3
        vsys                                 : vsys1
        application                          : bittorrent
        rule                                 : rule1
        session to be logged at end          : True
        session in session ager              : True
        session updated by HA peer           : False
        layer7 processing                    : completed
        URL filtering enabled                : False
        session via syn-cookies              : False
        session terminated on host           : False
        session traverses tunnel             : False
        captive portal session               : False
        ingress interface                    : ethernet1/21
        egress interface                     : ethernet1/22
        session QoS rule                     : N/A (class 4)
        tracker stage l7proc                 : ctd decoder bypass
        end-reason                           : unknown
```

Discard a Session Without a Commit

Perform this task to permanently discard a session, such as a session that is overloading the packet buffer. No commit is required; the session is discarded immediately after executing the command. The commands apply to both offloaded and non-offloaded sessions.

STEP 1 | In the CLI, execute the following operational command on any hardware model:

```
admin@PA-7050> request session-discard [timeout <seconds>] [reason <reason-string>] id <session-id>
```

The default timeout is 3,600 seconds.

STEP 2 | Verify that sessions have been discarded.

```
admin@PA-7050> show session all filter state discard
```


© 2018 Palo Alto Networks, Inc.

Certifications

The following topics describe how to configure Palo Alto Networks® firewalls and appliances to support the Common Criteria and the Federal Information Processing Standard 140-2 (FIPS 140-2), which are security certifications that ensure a standard set of security assurances and functionalities. These certifications are often required by civilian U.S. government agencies and government contractors.

For details about product certifications and third-party validation, refer to the Certifications page.

> Enable FIPS and Common Criteria Support
> FIPS-CC Security Functions

© 2018 Palo Alto Networks, Inc.

Enable FIPS and Common Criteria Support

Use the following procedures to enable FIPS-CC mode on a software version that supports Common Criteria and the Federal Information Processing Standards 140-2 (FIPS 140-2). When you enable FIPS-CC mode, all FIPS and CC functionality is included.

FIPS-CC mode is supported on all Palo Alto Networks next-generation firewalls and appliances—including VM-Series firewalls. To enable FIPS-CC mode, first boot the firewall into the Maintenance Recovery Tool (MRT) and then change the operational mode from normal mode to FIPS-CC mode. The procedure to change the operational mode is the same for all firewalls and appliances but the procedure to access the MRT varies.

When you enable FIPS-CC mode, the firewall will reset to the factory default settings; all configuration will be removed.

- Access the Maintenance Recovery Tool (MRT)
- Change the Operational Mode to FIPS-CC Mode

Access the Maintenance Recovery Tool (MRT)

The Maintenance Recovery Tool (MRT) enables you to perform several tasks on Palo Alto Networks firewalls and appliances. For example, you can revert the firewall or appliance to factory default settings, revert PAN-OS or a content update to a previous version, run diagnostics on the file system, gather system information, and extract logs. Additionally, you can use the MRT to Change the Operational Mode to FIPS-CC Mode or from FIPS-CC mode to normal mode.

The following procedures describe how to access the Maintenance Recovery Tool (MRT) on various Palo Alto Networks products.

- Access the MRT on hardware firewalls and appliances (such as PA-200 firewalls, PA-7000 Series firewalls, or M-Series appliances).
  1. Establish a serial console session to the firewall or appliance.
     1. Connect a serial cable from the serial port on your computer to the console port on the firewall or appliance.

        If your computer does not have a 9-pin serial port but does have a USB port, use a serial-to-USB converter to establish the connection. If the firewall has a micro USB console port, connect to the port using a standard Type-A USB to micro USB cable.
     2. Open terminal emulation software on your computer and set to 9600-8-N-1 and then connect to the appropriate COM port.

        On a Windows system, you can go to the Control Panel to view the COM port settings for Device and Printers to determine which COM port is assigned to the console.
     3. Log in using an administrator account. (The default username/password is admin/admin.)
  2. Enter the following CLI command and press **y** to confirm:

     ```
     debug system maintenance-mode
     ```
  3. After the firewall or appliance boots to the MRT welcome screen (in approximately 2 to 3 minutes), press Enter on **Continue** to access the MRT main menu.

© 2018 Palo Alto Networks, Inc.



You can also access the MRT by rebooting the firewall or appliance and entering **`maint`** *at the maintenance mode prompt. A direct serial console connection is required.*

After the firewall or appliance boots into the MRT, you can access the MRT remotely by establishing an SSH connection to the management (MGT) interface IP address. At the login prompt, enter **`maint`** as the username and the firewall or appliance serial number as the password.

- Access the MRT on VM-Series firewalls deployed in a private cloud (such as on a VMware ESXi or KVM hypervisor).
  1. Establish an SSH session to the management IP address of the firewall and log in using an administrator account.
  2. Enter the following CLI command and press **`y`** to confirm:

     ```
     debug system maintenance-mode
     ```

     It will take approximately 2 to 3 minutes for the firewall to boot to the MRT. During this time, your SSH session will disconnect.
  3. After the firewall boots to the MRT welcome screen, log in based on the operational mode:
     - **Normal mode**—Establish an SSH session to the management IP address of the firewall and log in using **`maint`** as the username and the firewall or appliance serial number as the password.
     - **FIPS-CC mode**—Access the virtual machine management utility (such as the vSphere client) and connect to the virtual machine console.
  4. From the MRT welcome screen, press Enter on **`Continue`** to access the MRT main menu.

- Access the MRT on VM-Series firewalls deployed in the public cloud (such as AWS or Azure).
  1. Establish an SSH session to the management IP address of the firewall and log in using an administrator account.
  2. Enter the following CLI command and press **`y`** to confirm:

     ```
     debug system maintenance-mode
     ```

     It will take approximately 2 to 3 minutes for the firewall to boot to the MRT. During this time, your SSH session will disconnect.
  3. After the firewall boots to the MRT welcome screen, log in based on the virtual machine type:
     - **AWS**—Log in as **`ec2-user`** and select the SSH public key associated with the virtual machine when you deployed it.
     - **Azure**—Enter the credentials you created when you deployed the VM-Series firewall.
  4. From the MRT welcome screen, press Enter on **`Continue`** to access the MRT main menu.

Change the Operational Mode to FIPS-CC Mode

The following procedure describes how to change the operational mode of a Palo Alto Networks product from normal mode to FIPS-CC mode.

STEP 1 | Connect to the firewall or appliance and Access the Maintenance Recovery Tool (MRT).

STEP 2 | Select **`Set FIPS-CC Mode`** from the menu.

© 2018 Palo Alto Networks, Inc.

STEP 3 | Select **Enable FIPS-CC Mode**. The mode change operation starts and a status indicator shows progress. After the mode change is complete, the status shows **Success**.

STEP 4 | When prompted, select **Reboot**.

> *If you change the operational mode on a VM-Series firewall deployed in the public cloud (AWS or Azure) and you lose your SSH connection to the MRT before you are able to* ***Reboot****, you must wait 10-15 minutes for the mode change to complete, log back into the MRT, and then reboot the firewall to complete the operation.*

After you switch to FIPS-CC mode, you see the following status: `FIPS-CC mode enabled successfully`.

In addition, the following changes are in effect:

- FIPS-CC displays at all times in the status bar at the bottom of the web interface.
- The default administrator login credentials change to admin/paloalto.

See FIPS-CC Security Functions for details on the security functions that are enforced in FIPS-CC mode.

© 2018 Palo Alto Networks, Inc.

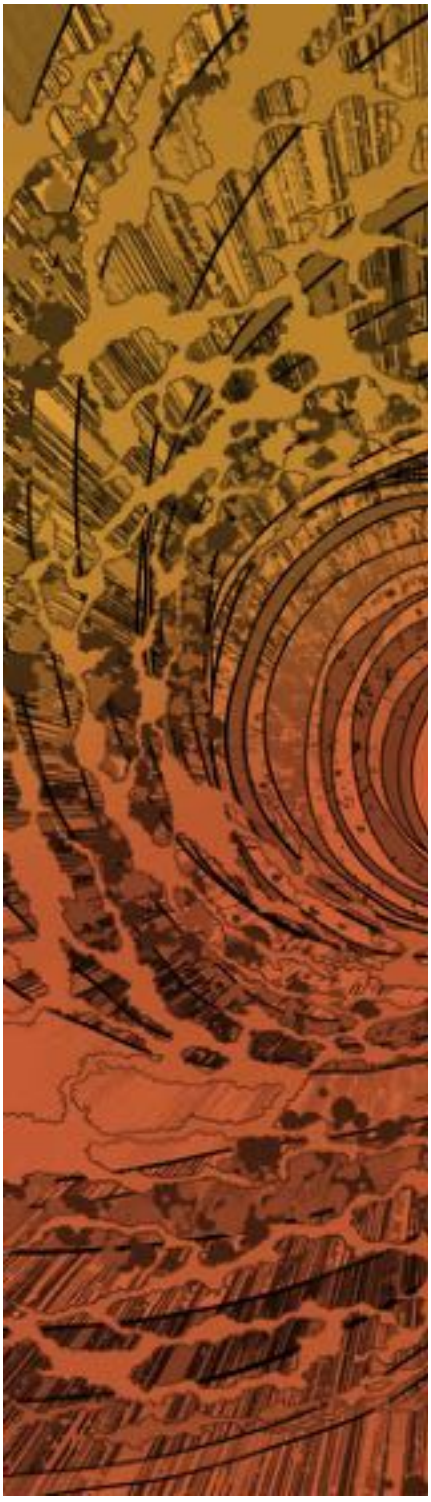
FIPS-CC Security Functions

When FIPS-CC mode is enabled, the following security functions are enforced on all firewalls and appliances:

- To log in, the browser must be TLS 1.1 (or later) compatible; on a WF-500 appliance, you manage the appliance only through the CLI and you must connect using an SSHv2-compatible client application.
- All passwords must be at least six characters.
- You must ensure that **Failed Attempts** and **Lockout Time (min)** are greater than 0 in authentication settings. If an administrator reaches the **Failed Attempts** threshold, the administrator is locked out for the duration defined in the **Lockout Time (min)** field.
- You must ensure that the **Idle Timeout** is greater than 0 in authentication settings. If a login session is idle for more than the specified time, the administrator is automatically logged out.
- The firewall or appliance automatically determines the appropriate level of self-testing and enforces the appropriate level of strength in encryption algorithms and cipher suites.
- Unapproved FIPS-CC algorithms are not decrypted—they are ignored during decryption.
- When configuring an IPSec VPN, the administrator must select a cipher suite option presented to them during the IPSec setup.
- Self-generated and imported certificates must contain public keys that are either RSA 2,048 bits (or more) or ECDSA 256 bits (or more); you must also use a digest of SHA256 or greater.

  *You cannot use a hardware security module (HSM) (*Secure Keys with a Hardware Security Module*) to store the private ECDSA keys used for* SSL Forward Proxy *or* SSL Inbound Inspection*.*
- Telnet, TFTP, and HTTP management connections are not available.
- High availability (HA) port encryption is required.
- The serial console port in FIPS-CC mode functions as a limited status output port only; CLI access is not available.
- The serial console port on hardware and private-cloud VM-Series firewalls booted into the MRT provides interactive access to the MRT.
- Interactive console access is not supported in the hypervisor environment private-cloud VM-Series firewalls booted into the MRT; you can access the MRT only using SSH.



paloalto NETWORKS®

Panorama Administrator's Guide

Version 7.0

Contact Information

Corporate Headquarters:

Palo Alto Networks

4401 Great America Parkway

Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-us

About this Guide

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

If you have an M-Series appliance, this guide takes over after you finish rack mounting your M-Series appliance.

For more information, refer to the following sources:

- For information on how to configure other components in the Palo Alto Networks Next-Generation Security Platform, go to the Technical Documentation portal: https://www.paloaltonetworks.com/documentation or search the documentation.
- For access to the knowledge base, complete documentation set, discussion forums, and videos, refer to https://live.paloaltonetworks.com.
- For contacting support, for information on support programs, to manage your account or devices, or to open a support case, refer to https://www.paloaltonetworks.com/support/tabs/overview.html.
- For the most current PAN-OS and Panorama 7.0 release notes, go to https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os-release-notes.html.

To provide feedback on the documentation, please write to us at: documentation@paloaltonetworks.com.

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2016 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at http://www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

Revision Date: September 30, 2016

Table of Contents

Panorama Overview 9
About Panorama 10
Panorama Platforms 11
Centralized Configuration and Deployment Management 13
Context Switch—Firewall or Panorama 13
Templates and Template Stacks 13
Device Groups 15
Centralized Logging and Reporting 19
Logging Options 19
Managed Collectors and Collector Groups 20
Caveats for a Collector Group with Multiple Log Collectors 20
Centralized Reporting 22
Panorama Commit Operations 23
Role-Based Access Control 25
Administrative Roles 25
Authentication Profiles and Sequences 26
Access Domains 27
Administrative Authentication 27
Panorama Recommended Deployments 29
Panorama for Centralized Management and Reporting 29
Panorama in a Distributed Log Collection Deployment 30
Plan Your Deployment 31
Deploy Panorama: Task Overview 33

Set Up Panorama 35
Determine Panorama Log Storage Requirements 36
Set Up the Panorama Virtual Appliance 38
Setup Prerequisites for the Panorama Virtual Appliance 38
Install the Panorama Virtual Appliance 39
Perform Initial Configuration of the Panorama Virtual Appliance 42
Expand Log Storage Capacity on the Panorama Virtual Appliance 44
Increase CPUs and Memory on the Panorama Virtual Appliance 46
Complete the Panorama Virtual Appliance Setup 48
Set Up the M-Series Appliance 49
Perform Initial Configuration of the M-Series Appliance 50
Set up the M-Series Appliance as a Log Collector 53
Increase Storage on the M-Series Appliance 58
Register Panorama and Install Licenses 61
Register Panorama 61
Activate a Panorama Support License 62

© Palo Alto Networks, Inc.

Activate/Retrieve a Device Management License on the Panorama Virtual Appliance . . . 63
Activate/Retrieve a Device Management License on the M-Series Appliance 63
Install Content and Software Updates for Panorama 65
Panorama, Log Collector, and Firewall Version Compatibility 65
Install Updates for Panorama with HA Configuration 65
Install Updates for Panorama with Internet Connection 66
Install Updates for Panorama without Internet Connection 69
Transition to a Different Panorama Platform 73
Migrate from a Panorama Virtual Appliance to an M-Series Appliance 73
Migrate from an M-100 Appliance to an M-500 Appliance 76
Access and Navigate Panorama Management Interfaces 79
Log in to the Panorama Web Interface 79
Navigate the Panorama Web Interface 79
Log in to the Panorama CLI 80
Set Up Administrative Access to Panorama 82
Configure an Admin Role Profile 82
Configure an Access Domain 82
Configure Administrative Accounts and Authentication 83
Configure an Administrative Account 83
Configure an Administrator with Kerberos SSO, External, or Local Authentication 84
Configure an Administrator with Certificate-Based Authentication for the Web Interface . 85
Configure an Administrator with SSH Key-Based Authentication for the CLI 86
Configure RADIUS Vendor-Specific Attributes for Administrator Authentication 87

Manage Firewalls 89

Add a Firewall as a Managed Device 90
Manage Device Groups 91
Add a Device Group 91
Create a Device Group Hierarchy 92
Create Objects for Use in Shared or Device Group Policy 93
Revert to Inherited Object Values 95
Manage Unused Shared Objects 95
Manage Precedence of Inherited Objects 96
Move or Clone a Policy Rule or Object to a Different Device Group 96
Select a URL Filtering Vendor on Panorama 97
Push a Policy Rule to a Subset of Firewalls 101
Manage the Rule Hierarchy 102
Manage Templates and Template Stacks 104
Template Capabilities and Exceptions 104
Add a Template 104
Configure a Template Stack 106
Override a Template Setting 107
Disable/Remove Template Settings 108
Transition a Firewall to Panorama Management 109
Plan the Transition to Panorama Management 109
Migrate a Firewall to Panorama Management 110
Load a Partial Firewall Configuration into Panorama 113

 © Palo Alto Networks, Inc.

Use Case: Configure Firewalls Using Panorama . . . 116
Device Groups . . . 116
Templates . . . 117
Set Up Your Centralized Configuration and Policies . . . 118

Manage Log Collection . . . 125

Configure a Managed Collector . . . 126

Manage Collector Groups . . . 130
Configure a Collector Group . . . 130
Move a Log Collector to a Different Collector Group . . . 133
Remove a Firewall from a Collector Group . . . 135

Configure Log Forwarding to Panorama . . . 136

Verify Log Forwarding to Panorama . . . 138

Modify Log Forwarding and Buffering Defaults . . . 139

Configure Log Forwarding from Panorama to External Destinations . . . 141

Log Collection Deployments . . . 143
Plan a Log Collection Deployment . . . 143
Deploy Panorama with Dedicated Log Collectors . . . 145
Deploy Panorama with Default Log Collectors . . . 152
Deploy Panorama Virtual Appliances with Local Log Collection . . . 158

Manage Licenses and Updates . . . 161

Manage Licenses on Firewalls Using Panorama . . . 162

Deploy Updates to Firewalls and Log Collectors Using Panorama . . . 163
Supported Updates . . . 163
Schedule a Content Update Using Panorama . . . 164
Deploy an Update to Log Collectors when Panorama is Internet-connected . . . 164
Deploy an Update to Log Collectors when Panorama is not Internet-connected . . . 166
Deploy an Update to Firewalls when Panorama is Internet-connected . . . 168
Deploy an Update to Firewalls when Panorama is not Internet-connected . . . 169

Monitor Network Activity . . . 173

Use Panorama for Visibility . . . 174
Monitor the Network with the ACC and AppScope . . . 174
Analyze Log Data . . . 176
Generate, Schedule, and Email Reports . . . 177

Use Case: Monitor Applications Using Panorama . . . 180

Use Case: Respond to an Incident Using Panorama . . . 184
Incident Notification . . . 184
Review the Widgets in the ACC . . . 185
Review Threat Logs . . . 186
Review WildFire Logs . . . 187
Review Data Filtering Logs . . . 188
Update Security Rules . . . 188

Panorama High Availability 191
Panorama HA Prerequisites 192
Priority and Failover on Panorama in HA 193
Failover Triggers 195
HA Heartbeat Polling and Hello Messages 195
HA Path Monitoring 195
Logging Considerations in Panorama HA 197
Logging Failover on a Panorama Virtual Appliance 197
Logging Failover on an M-Series Appliance 197
Synchronization Between Panorama HA Peers 199
Manage a Panorama HA Pair 200
Set Up HA on Panorama 200
Test Panorama HA Failover 202
Switch Priority after Panorama Failover to Resume NFS Logging 202
Restore the Primary Panorama to the Active State 203

Administer Panorama 205
Manage Configuration Backups 206
Schedule Export of Configuration Files 206
Manage Panorama and Firewall Configuration Backups 207
Configure the Number of Configuration Backups Panorama Stores 208
Load a Configuration Backup on a Managed Firewall 209
Compare Changes in Panorama Configurations 210
Validate a Panorama Configuration 211
Restrict Access to Configuration Changes 212
Types of Locks 212
Locations for Taking a Lock 212
Take a Lock 213
View Lock Holders 213
Enable Automatic Acquisition of the Commit Lock 213
Remove a Lock 214
Add Custom Logos to Panorama 215
View Panorama Task Completion History 216
Manage Storage Quotas and Expiration Periods for Logs and Reports 217
Log and Report Storage 217
Log and Report Expiration Periods 217
Configure Storage Quotas and Expiration Periods for Logs and Reports 218
Monitor Panorama 220
Panorama System and Configuration Logs 220
Monitor Panorama and Log Collector Statistics Using SNMP 221
Reboot or Shut Down Panorama 223
Configure Panorama Password Profiles and Complexity 224

 © Palo Alto Networks, Inc.

Troubleshooting 227
Troubleshoot Panorama System Issues 228
Generate Diagnostic Files for Panorama 228
Diagnose Panorama Suspended State 228
Monitor the File System Integrity Check 228
Manage Panorama Storage for Software and Content Updates 229
Recover from Split Brain in Panorama HA Deployments 229
Troubleshoot Log Storage and Connection Issues. 231
Verify Panorama Port Usage 231
Resolve Zero Log Storage for a Collector Group 232
Replace a Failed Disk on an M-Series Appliance 232
Replace the Virtual Disk on a Panorama ESXi Server 233
Replace the Virtual Disk on vCloud Air 233
Migrate Logs to a New M-Series Appliance in Log Collector Mode 234
Migrate Logs to a New M-Series Appliance in Panorama Mode 238
Recover Logs after Panorama Failure/RMA in Non-HA Deployments. 243
Regenerate Metadata for M-Series Appliance RAID Pairs 245
Replace an RMA Firewall 247
Partial Device State Generation for Firewalls 247
Before Starting RMA Firewall Replacement 247
Restore the Firewall Configuration after Replacement 249
Troubleshoot Commit Failures 251
Troubleshoot Registration or Serial Number Errors 252
Troubleshoot Reporting Errors 253
View Task Success or Failure Status 254

© Palo Alto Networks, Inc.

© Palo Alto Networks, Inc.

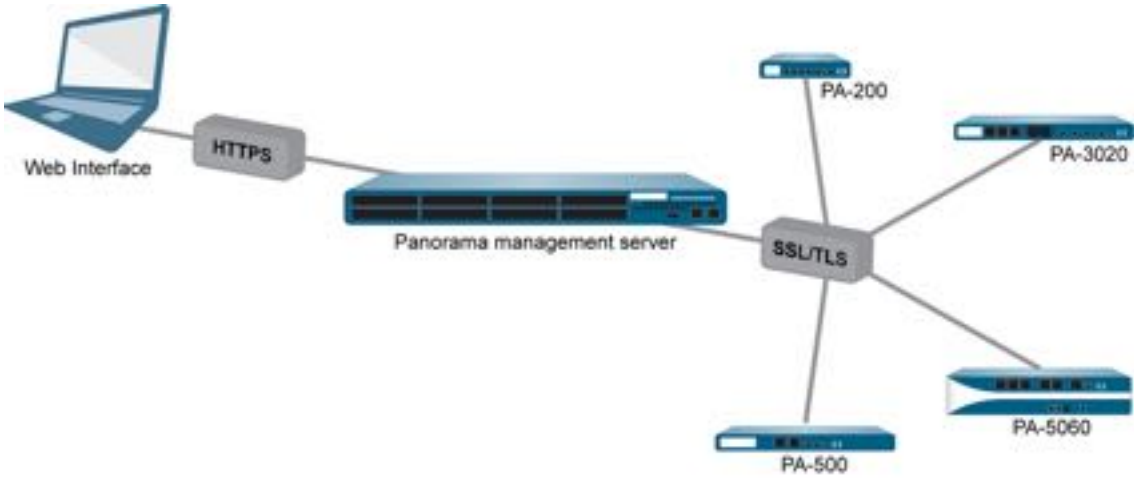
Panorama Overview

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

- About Panorama
- Panorama Platforms
- Centralized Configuration and Deployment Management
- Centralized Logging and Reporting
- Panorama Commit Operations
- Role-Based Access Control
- Panorama Recommended Deployments
- Plan Your Deployment
- Deploy Panorama: Task Overview

About Panorama

Panorama provides centralized management of Palo Alto Networks next-generation firewalls, as the following figure illustrates:



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

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

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

 © Palo Alto Networks, Inc.

Panorama Platforms

Panorama is available in the following platforms, each of which supports licenses for managing up to 25, 100, or 1,000 firewalls:

- **Panorama virtual appliance**—You can install the Panorama virtual appliance on a VMware ESXi server or in VMware vCloud Air. The virtual appliance allows for a simple installation and facilitates server consolidation for sites that need a virtual management appliance. It also supports integration with a Network File Share (NFS) system for increased storage and log retention beyond 2TB. The Panorama virtual appliance works best in environments with logging rates of up to 10,000 logs per second. You can forward firewall logs directly to the Panorama virtual appliance (see Deploy Panorama Virtual Appliances with Local Log Collection) or use the Panorama virtual appliance to manage Dedicated Log Collectors that are M-Series appliances (see Deploy Panorama with Dedicated Log Collectors).
- **M-Series appliance**—The M-100 appliance and M-500 appliance are dedicated hardware platforms intended for large-scale deployments. In environments with high logging rates (over 10,000 logs per second) and log retention requirements, these appliances enables scaling of your log collection infrastructure. Both appliances use RAID drives to store firewall logs and support RAID 1 mirroring to protect against disk failures. Both appliances use an SSD to store the logs that Panorama and Log Collectors generate. Only the M-500 appliance has redundant, hot-swappable power supplies and front-to-back airflow. The M-500 appliance also has faster processors and greater memory for better performance (for example, faster commit times). These attributes make the M-500 appliance more suitable for datacenters than the M-100 appliance. The log storage capacity and maximum log collection rate varies by appliance:

Appliance	SSD Storage	Default RAID Storage	Maximum RAID Storage	Maximum Logging Rate
M-100 appliance	120GB	2 drives (1TB total)	8 drives (4TB total)	30,000 logs/second
M-500 appliance	240GB	8 drives (4TB total)	16 drives (8TB total)	60,000 logs/second

  You can deploy the M-Series appliance in the following modes to separate the central management function from the log collection function:
  - **Panorama mode**: The appliance performs both central management and log collection. This is the default mode. For configuration details, see Deploy Panorama with Default Log Collectors.
  - **Log Collector mode**: The appliance functions as a Dedicated Log Collector. If multiple firewalls forward large volumes of log data, the M-Series appliance in Log Collector mode provides increased scale and performance. In this mode, the appliance has no web interface for administrative access, only a command line interface (CLI). However, you can manage the appliance using the web interface of the Panorama management server (M-Series appliance in Panorama mode or a Panorama virtual appliance). CLI access to an M-Series appliance in Log Collector mode is only necessary for initial setup and debugging. For configuration details, see Deploy Panorama with Dedicated Log Collectors.

  For more details and specifications, see the M-100 and M-500 Hardware Reference Guides.

The platform choice depends on your need for a virtual appliance and your log collection requirements (see Determine Panorama Log Storage Requirements):

Log Collection Rate	Platform
Up to 10,000 logs/second	Panorama virtual appliance
Up to 30,000 logs/second	M-100 appliance
Up to 60,000 logs/second	M-500 appliance

© Palo Alto Networks, Inc.

Centralized Configuration and Deployment Management

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

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

Context Switch—Firewall or Panorama

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

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

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

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

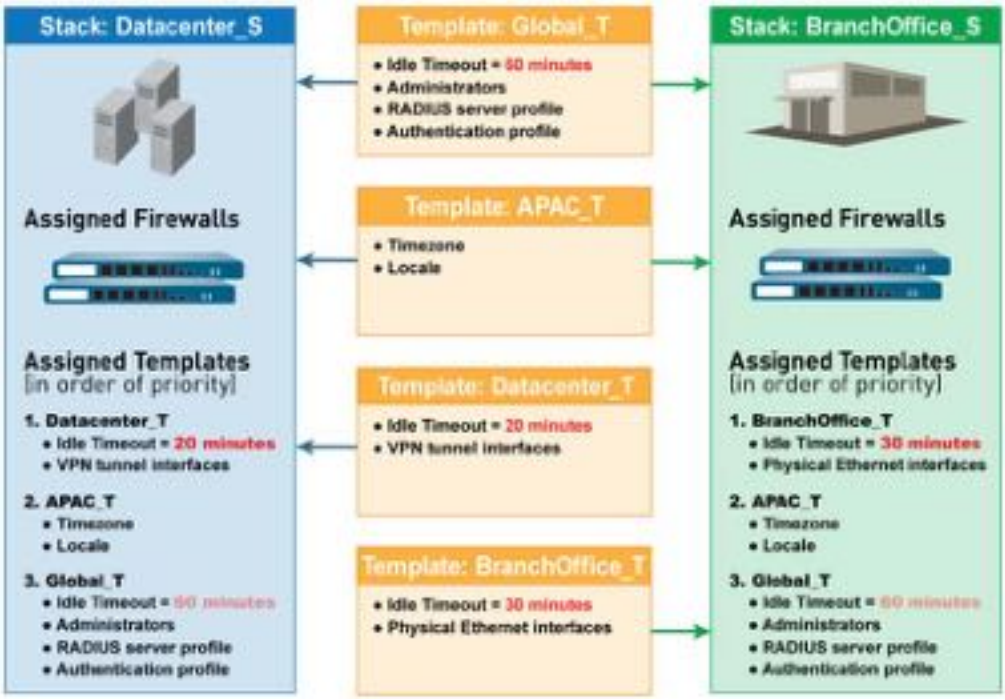
Templates and Template Stacks

You use templates to configure the settings that enable firewalls to operate on the network. Templates enable you to define a common base configuration using the **Network** and **Device** tabs on Panorama. For example, you can use templates to manage interface and zone configurations, server profiles for logging and syslog access, and network profiles for controlling access to zones and IKE gateways. When defining a template, consider assigning firewalls that are the same hardware model and require access to similar network resources, such as gateways and syslog servers.

If your network has groups of firewalls with some group-specific settings and some settings that are common across groups, you can simplify management by assigning the firewalls to a template stack for each group. A template stack is a combination of templates: the assigned firewalls inherit the settings from every template in the stack. This enables you to avoid the redundancy of adding every setting to every template. The following figure illustrates an example deployment in which you assign data center firewalls in the

Asia-Pacific (APAC) region to a stack that has one template with global settings, one template with APAC-specific settings, and one template with data center-specific settings. To manage firewalls in an APAC branch office, you can then re-use the global and APAC-specific templates by adding them to another stack that includes a template with branch-specific settings. Templates in a stack have a configurable priority order that ensures Panorama pushes only one value for any duplicate setting. Panorama evaluates the templates listed in a stack configuration from top to bottom, with higher templates having priority. The following figure illustrates a data center stack in which the data center template has a higher priority than the global template: Panorama pushes the idle timeout value from the data center template and ignores the value from the global template.

Figure: Template Stacks



To accommodate firewalls that have unique settings, you can use templates (single or stacked) to push a limited common base configuration to all firewalls, and in individual firewalls configure device-specific settings. Alternatively, you can push a broader common base configuration and in the individual firewalls override certain pushed settings with device-specific values. When you override a setting, the firewall saves that setting to its local configuration; Panorama no longer manages the setting. To restore template values after overriding them, you can use Panorama to force the template configuration onto a firewall. For example, after defining a common NTP server in a template and overriding the NTP server configuration on a firewall to accommodate its local time zone, you can later revert to the NTP server defined in the template.

You cannot use templates to set firewall modes: virtual private network (VPN) mode, multiple virtual systems mode (multi-vsys mode), and operational mode (normal, Federal Information Processing Standards [FIPS], or Common Criteria [CC]). For details, see Template Capabilities and Exceptions. However, you can assign firewalls that have non-matching modes to the same template or stack. In such cases, Panorama pushes mode-specific settings only to firewalls that support those modes. As an exception, you can configure Panorama to push the settings of the default vsys in a template to firewalls that don't support virtual systems or have none configured.

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

 © Palo Alto Networks, Inc.

Device Groups

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

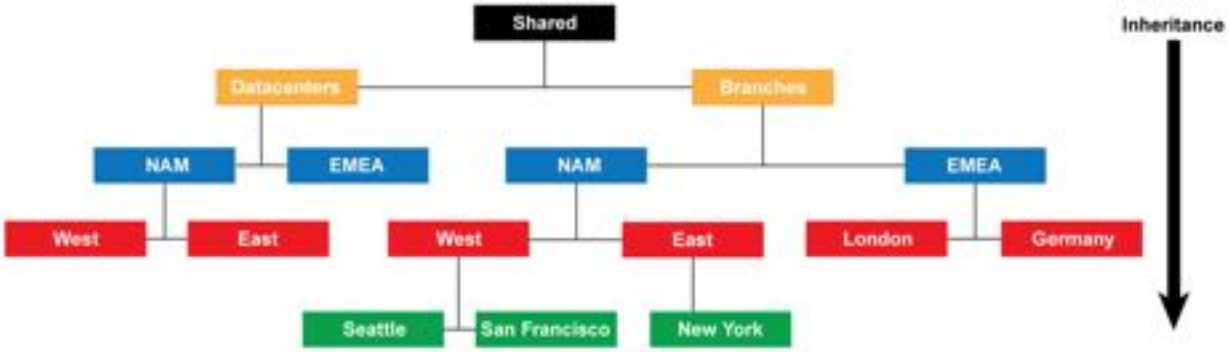
- Device Group Hierarchy
- Device Group Policies
- Device Group Objects

Device Group Hierarchy

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

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

Figure: Device Group Hierarchy



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

Device Group Policies

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

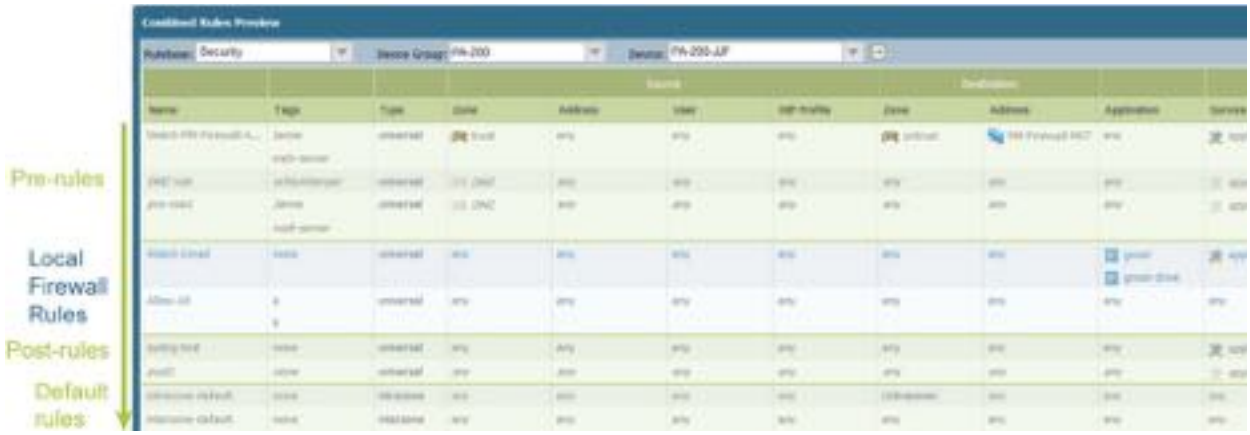
Evaluation Order	Rule Scope and Description	Administration Platform
Shared pre-rules Device group pre-rules	Panorama pushes shared pre-rules to all the firewalls in all device groups. Panorama pushes device group-specific pre-rules to all the firewalls in a particular device group and its descendant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

 © Palo Alto Networks, Inc.

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

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 appear in green, while local firewall rules appear in blue between the pre-rules and post-rules.

Figure: Rule Hierarchy



Device Group Objects

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

Create Objects for Use in Shared or Device Group Policy once and use them many times, you reduce administrative overhead and ensure consistency across firewall policies.

You can configure how Panorama handles objects system-wide:

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

 © Palo Alto Networks, Inc.

Centralized Logging and Reporting

Panorama aggregates data from all managed firewalls and provides visibility across all the traffic on the network. It also provides an audit trail for all policy modifications and configuration changes made to the managed firewalls. In addition to aggregating logs, Panorama can aggregate and forward Simple Network Management Protocol (SNMP) traps, email notifications, and syslog messages to an external destination.

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

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

- Logging Options
- Managed Collectors and Collector Groups
- Caveats for a Collector Group with Multiple Log Collectors
- Centralized Reporting

Logging Options

Both the Panorama virtual appliance and M-Series appliance can collect logs that the managed firewalls forward. You can then Configure Log Forwarding from Panorama to External Destinations (syslog server, email server, or Simple Network Management Protocol [SNMP] trap server). The logging options vary on each Panorama platform.



The PA-7000 Series firewall can't forward logs to Panorama, only to external services directly. However, when you monitor logs or generate reports for a device group that includes a PA-7000 Series firewall, Panorama queries the firewall in real-time to display its log data.

Panorama Platform	Logging Options
Virtual appliance	Offers three logging options: • Use the approximately 11GB of internal storage space allocated for logging as soon as you install the virtual appliance. • Add a virtual disk that can support up to 2TB of storage. • Mount a Network File System (NFS) datastore in which you can configure the storage capacity that is allocated for logging.

Panorama Platform	Logging Options
M-Series appliance	The default shipping configuration for the M-100 appliance includes two disks with a total of 1TB storage capacity. For the M-500 appliance, the default configuration includes eight disks for 4TB of storage. Both appliances use RAID 1 to protect against disk failures. You can Increase Storage on the M-Series Appliance to 4TB on the M-100 appliance and 8TB on the M-500 appliance. When an M-Series appliance is in Panorama mode, you can enable the RAID disks to serve as the default Log Collector. If you have an M-Series appliance is in Log Collector mode (Dedicated Log Collector), you use Panorama to assign firewalls to the Dedicated Log Collectors. In a deployment with multiple Dedicated Log Collectors, Panorama queries all managed Log Collectors to generate an aggregated view of traffic and cohesive reports. For easy scaling, begin with a single Panorama and incrementally add Dedicated Log Collectors as your needs expand.

Managed Collectors and Collector Groups

A Log Collector can be local to an M-Series appliance in *Panorama mode* (default Log Collector) or can be an M-Series appliance in *Log Collector mode* (Dedicated Log Collector). Because you use Panorama to configure and manage Log Collectors, they are also known as *managed collectors*. An M-Series appliance in Panorama mode or a Panorama virtual appliance can manage Dedicated Log Collectors. To administer Dedicated Log Collectors using the Panorama web interface, you must add them as managed collectors. Otherwise, administrative access to a Dedicated Log Collector is only available through its CLI using the default administrative user (*admin*) account. Dedicated Log Collectors do not support additional administrative user accounts.

A Collector Group is one or more managed collectors that operate as a single logical log collection unit. If the group contains Dedicated Log Collectors, the logs are uniformly distributed across all the disks in each Log Collector and across all members in the Collector Group. This distribution maximizes the use of the available storage space. To manage a Log Collector, you must add it to a Collector Group. If you assign more than one Log Collector to a Collector Group, see Caveats for a Collector Group with Multiple Log Collectors.

The Collector Group configuration specifies which managed firewalls can send logs to the Log Collectors in the group. After you configure the Log Collectors and enable the firewalls to forward logs, each firewall forwards its logs to the assigned Log Collector.

Managed collectors and Collector Groups are integral to a distributed log collection deployment on Panorama. A distributed log collection deployment allows for easy scalability and incremental addition of Dedicated Log Collectors as your logging needs grow. The M-Series appliance in Panorama mode can log to its default Collector Group and then be expanded to a distributed log collection deployment with one or more Collector Groups that include Dedicated Log Collectors.

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

Caveats for a Collector Group with Multiple Log Collectors

You can Configure a Collector Group with multiple Log Collectors to ensure log redundancy or to accommodate logging rates that exceed the capacity of a single Log Collector (see Panorama Platforms). All the Log Collectors in any particular Collector Group must be the same platform: all M-500 appliances or all

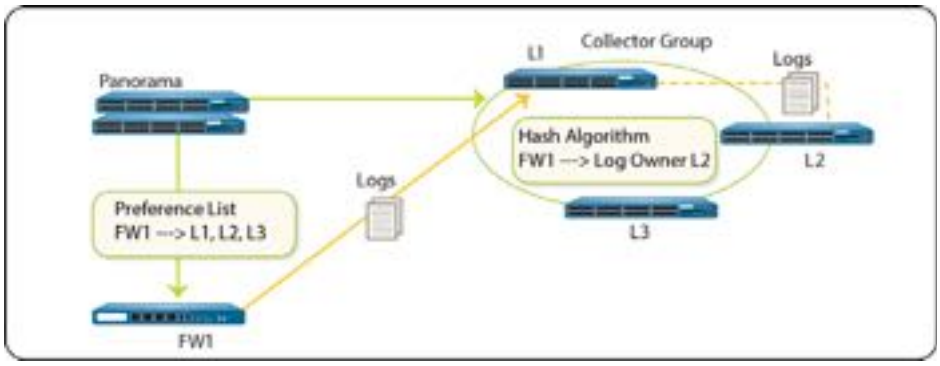
 © Palo Alto Networks, Inc.

M-100 appliances. For example, if a single managed firewall generates 16TB of logs, the Collector Group that receives those logs will require at least four Log Collectors that are M-100 appliances or two Log Collectors that are M-500 appliances.

A Collector Group with multiple Log Collectors uses the available storage space as one logical unit and uniformly distributes the logs across all its Log Collectors. The log distribution is based on the disk capacity of the Log Collectors (1TB to 8TB, depending on the number of disk pairs and the M-Series platform) and a hash algorithm that dynamically decides which Log Collector owns the logs and writes to disk. Although Panorama uses a preference list to prioritize the list of Log Collectors to which a managed firewall can forward logs, Panorama does not necessarily write the logs to the first Log Collector specified in the preference list. For example, consider the following preference list:

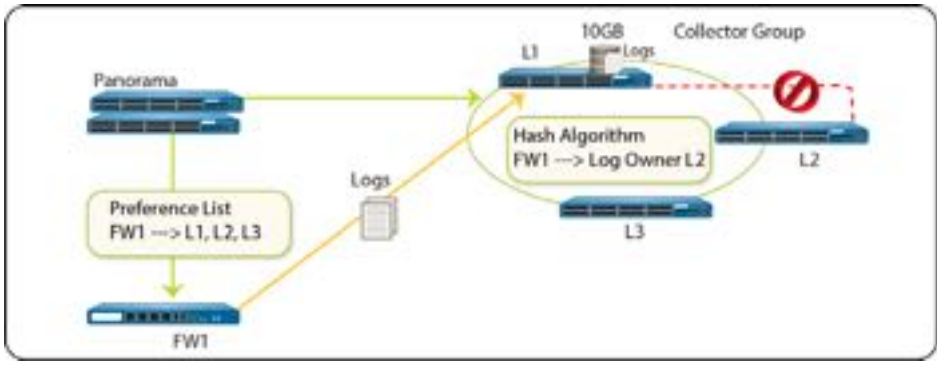
Managed Firewall	Log Forwarding Preference List Defined on a Collector Group
FW1	L1,L2,L3
FW2	L4,L5,L6

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



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

With multiple Log Collectors in a Collector Group, the firewall does not buffer logs to its local storage when it can connect to its primary Log Collector. Therefore, FW1 will continue sending logs to L1. Because L2 is unavailable, the primary Log Collector L1 buffers the logs to its HDD, which has 10GB of log space. If L2 remains unavailable and the logs pending for L2 exceed 10GB, L1 will overwrite the older log entries to continue logging. In such an event, loss of logs is a risk.



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

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

  Because enabling redundancy creates more logs, this configuration requires more storage capacity. When a Collector Group runs out of space, it deletes older logs.
  Enabling redundancy doubles the log processing traffic in a Collector Group, which reduces its maximum logging rate by half, as each Log Collector must distribute a copy of each log it receives.

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

Centralized Reporting

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

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

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

For more information, see Monitor Network Activity.

 © Palo Alto Networks, Inc.

Panorama Commit Operations

When editing the configuration on Panorama, you are changing the candidate configuration file. The candidate configuration is a copy of the running configuration along with any changes you made since the last commit. The Panorama web interface displays all the configuration changes immediately. However, Panorama won't implement the changes until you commit them. The commit process validates the changes in the candidate configuration file and saves it as the running configuration on Panorama.



After any system event or administrator action causes Panorama to reboot, all your changes since the last commit will be lost. To preserve changes without committing them, periodically click **Save** at the top right of the web interface to save a snapshot of the candidate configuration. If a reboot occurs, you can then revert to the snapshot. For details on backing up and restoring running and candidate configurations, see Manage Panorama and Firewall Configuration Backups.

When initiating a commit on Panorama, select one of the following types:

Commit Options	Description
Panorama	Commits the changes on the current candidate configuration to the running configuration on Panorama. You must first commit your changes on Panorama, before committing any configuration updates (templates or device groups) to the managed firewalls or Collector Groups.
Template	Commits network and device configurations from a Panorama template or template stack to the selected firewalls.
Device Group	Commits policies and objects configured from Panorama to the selected firewalls/virtual systems.
Collector Group	Commits changes to the specified Collector Groups that Panorama manages.

When you perform a commit, Panorama pushes the entire configuration to the managed firewalls or Log Collectors. When the commit completes, a result displays: **Commit succeeded** or **Commit succeeded with warnings**.



Panorama can't perform a device group or template commit to firewalls while a local commit is in progress on those firewalls. The local commit can be manual (you click **Commit**) or automatic. PAN-OS performs an automatic commit when you downgrade content versions (for example, the WildFire version), or refresh address objects, FQDNs, or dynamic block lists.

Some other commit choices are:

- **Preview Changes**—This option is available when the **Commit Type** is **Panorama**. It enables you to compare the candidate configuration with the running configuration in the same way as the **Panorama > Config Audit** feature (see Compare Changes in Panorama Configurations). After clicking **Preview Changes**, select the number of lines to include for context, and click **OK**. As a best practice, preview your configuration changes before committing them.



Because the preview results display in a new window, your browser must allow pop-up windows. If the preview window does not open, refer to your browser documentation for the steps to unblock pop-up windows.

- **Validate Changes**—This option is available when the **Commit Type** is **Panorama**, **Template**, or **Device Group**. It enables you to Validate a Panorama Configuration before committing it.

- **Include Device and Network Templates**—This option is available when committing a device group from Panorama. It allows you to commit both device group and template or template stack changes to the selected firewalls in a single commit operation. The check box is selected by default. If you prefer to commit these changes as separate operations, clear this check box.
- **Force Template Values**—When performing a template or device group commit, the **Force Template Values** option overrides all local configuration and removes objects on the selected firewalls or virtual systems that don't exist in the template or template stack, or are overridden in the local configuration. This is an override that reverts all existing configuration on the managed firewall, and ensures that the firewall inherits only the settings defined in the template or template stack.
- **Merge with Candidate Config**—When enabled, this option allows you to merge and commit the Panorama configuration changes with any pending configuration changes that were implemented locally on the target firewall. If this option is not enabled, the candidate configuration on the firewall is not included in the commit operation. As a best practice, leave this option disabled if you allow firewall administrators to modify the configuration directly on a firewall and you don't want to include their changes when committing changes from Panorama.

 © Palo Alto Networks, Inc.

Role-Based Access Control

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



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

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

Administrative Roles

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

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

Dynamic Role	Privileges
Superuser	Full read-write access to Panorama
Superuser (read-only)	Read-only access to Panorama

Dynamic Role	Privileges
Panorama administrator	Full access to Panorama except for the following actions: • Create, modify, or delete Panorama or device administrators and roles. • Export, validate, revert, save, load, or import a configuration in the **Device > Setup > Operations** page. • Configure **Scheduled Config Export** functionality in the **Panorama** tab.

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

Admin Role Profile	Description
Panorama	For these roles, you can assign read-write access, read-only access, or no access to all the Panorama features that are available to the superuser dynamic role except the management of Panorama administrators and Panorama roles. For the latter two features, you can assign read-only access or no access, but you cannot assign read-write access. An example use of a Panorama role would be for security administrators who require access to security policy definitions, logs, and reports on Panorama.

Device Group and Template	For these roles, you can assign read-write access, read-only access, or no access to specific functional areas within device groups, templates, and firewall contexts. By combining these roles with Access Domains, you can enforce the separation of information among the functional or regional areas of your organization. Device Group and Template roles have the following limitations:
	• No access to the CLI or XML API
	• No access to configuration or system logs
	• No access to VM information sources
	• In the **Panorama** tab, access is limited to:
	  • Device deployment features (read-write, read-only, or no access)
	  • The device groups specified in the administrator account (read-write, read-only, or no access)
	  • The templates and managed devices specified in the administrator account (read-only or no access)
	An example use of this role would be for administrators in your operations staff who require access to the device and network configuration areas of the web interface for specific device groups and/or templates.

Authentication Profiles and Sequences

An authentication profile specifies the authentication service that validates the credentials of an administrator during login and defines how Panorama accesses the service. If you create a local administrator account on Panorama, you can authenticate the administrator to the local database, use an external service (RADIUS, TACACS+, LDAP, or Kerberos server), or use Kerberos single sign-on (SSO). If you use an external service, you must configure a server profile before you Configure an Admin Role Profile. If you want to use an external service for both account administration (instead of creating local accounts) and for authentication, you must Configure RADIUS Vendor-Specific Attributes for Administrator Authentication.

 © Palo Alto Networks, Inc.

Some environments have multiple databases for different users and user groups. To authenticate to multiple authentication sources (for example, local database and LDAP), configure an authentication sequence. An authentication sequence is a ranked order of authentication profiles that an administrator is matched against when logging in. Panorama checks against the local database first, and then checks each profile in sequence until the administrator is successfully authenticated. The administrator is denied access to Panorama only if authentication fails for all the profiles defined in the authentication sequence.

Access Domains

Access domains control administrative access to specific device groups (to manage policies and objects) and templates (to manage network and device settings), and also control the ability to switch context to the web interface of managed firewalls. Access domains apply only to administrators with Device Group and Template roles. By combining access domains with Administrative Roles, you can enforce the separation of information among the functional or regional areas of your organization.

You can manage access domains locally or by using RADIUS Vendor-Specific Attributes (VSAs). To use RADIUS VSAs, your network requires an existing RADIUS server and you must configure a RADIUS server profile to define how Panorama accesses the server. On the RADIUS server, you define a VSA attribute number and value for each administrator. The value defined must match the access domain configured on Panorama. When an administrator tries to log in to Panorama, Panorama queries the RADIUS server for the administrator access domain and attribute number. Based on the response from the RADIUS server, the administrator is authorized for access and is restricted to the firewalls, virtual systems, device groups, and templates that are assigned to the access domain.

For the relevant procedures, see:

- Configure an Access Domain.
- Configure RADIUS Vendor-Specific Attributes for Administrator Authentication.

Administrative Authentication

The following methods are available to authenticate Panorama administrators:

- **Local administrator account with local authentication**—Both the administrator account credentials and the authentication mechanisms are local to Panorama. To further secure the local administrator account, create a password profile that defines a validity period for passwords and set Panorama-wide password complexity settings. For details on how to configure this type of administrative access, see Configure an Administrator with Kerberos SSO, External, or Local Authentication.
- **Local administrator account with certificate- or key-based authentication**—With this option, the administrator accounts are local to Panorama, but authentication is based on Secure Shell (SSH) keys (for CLI access) or client certificates/common access cards (for the web interface). For details on how to configure this type of administrative access, see Configure an Administrator with Certificate-Based Authentication for the Web Interface and Configure an Administrator with SSH Key-Based Authentication for the CLI.
- **Local administrator account with external authentication**—The administrator accounts are managed on Panorama, but existing external authentication services (LDAP, Kerberos, TACACS+, or RADIUS) handle the authentication functions. If your network supports Kerberos single sign-on (SSO), you can configure

external authentication as an alternative in case SSO fails. For details on how to configure this type of administrative access, see Configure an Administrator with Kerberos SSO, External, or Local Authentication.

- **External administrator account and authentication**—An external RADIUS server handles account administration and authentication. To use this option, you must define Vendor-Specific Attributes (VSAs) on your RADIUS server that map to the administrator roles and access domains. For a high-level overview of the process, see Configure RADIUS Vendor-Specific Attributes for Administrator Authentication. For details on how to configure this type of administrative access, refer to Radius Vendor-Specific Attributes (VSAs).

 © Palo Alto Networks, Inc.

Panorama Recommended Deployments

A Panorama deployment comprises the Panorama management server (which has a browser-based interface), optional Log Collectors, and the Palo Alto Networks firewalls that Panorama manages. The recommended deployments are:

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

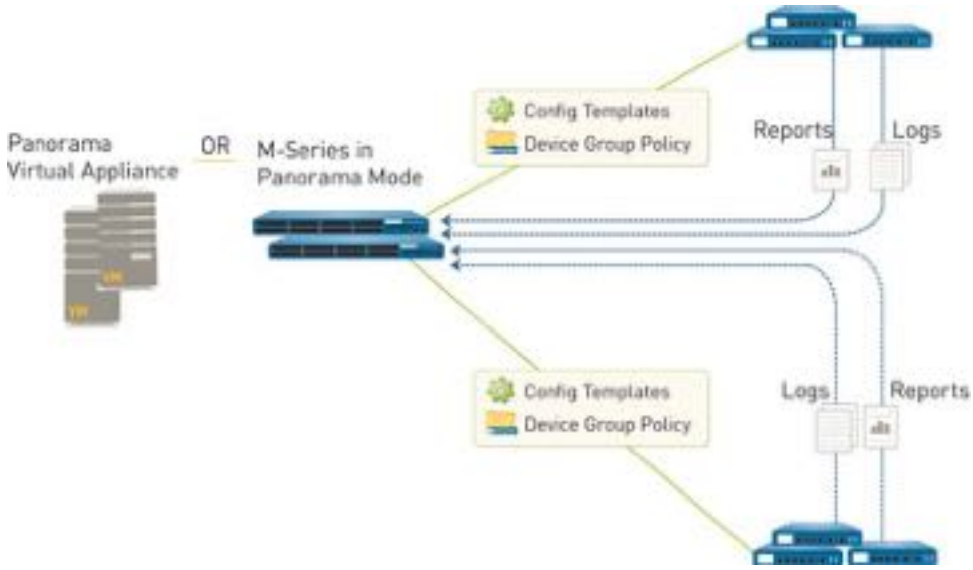


For the procedures to configure the most typical log collection deployments, see Log Collection Deployments.

Panorama for Centralized Management and Reporting

The following diagram illustrates how you can deploy the Panorama virtual appliance or M-Series appliance in a redundant configuration for the following benefits:

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



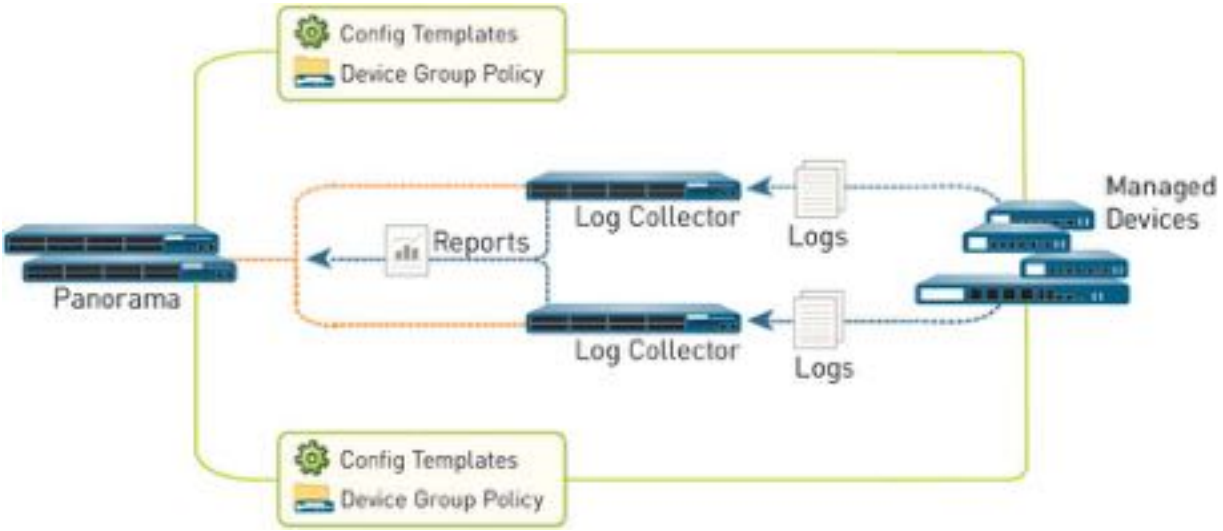
Panorama in a Distributed Log Collection Deployment

You can deploy the hardware-based Panorama—the M-Series appliance—either as a Panorama management server that performs management and log collection functions or as a Dedicated Log Collector that provides a comprehensive log collection solution for the firewalls on your network. Using the M-Series appliance as a Log Collector allows for a more robust environment where the log collection process is offloaded to a dedicated appliance. Using a dedicated appliance in a distributed log collection (DLC) deployment provides redundancy, improved scalability, and capacity for longer term log storage.

In a DLC deployment, the Panorama management server (Panorama virtual appliance or an M-Series appliance in Panorama mode) manages the firewalls and the Log Collectors. Using Panorama, you configure the firewalls to send logs to one or more Log Collectors. You can then use Panorama to query the Log Collectors and provide an aggregated view of network traffic. In a DLC configuration, you can access the logs stored on the Log Collectors from both the primary and secondary Panorama peers in a high availability (HA) pair.

In the following topology, the Panorama peers in an HA configuration manage the deployment and configuration of firewalls. This solution provides the following benefits:

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



 © Palo Alto Networks, Inc.

Plan Your Deployment

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

  If you already deployed and configured the Palo Alto Networks firewalls on your network, determine whether to transition the devices to centralized management. This process requires a migration of all configuration and policies from your firewalls to Panorama. For details, see Transition a Firewall to Panorama Management.
- Verify that Panorama is on the same release version or a later version than the firewalls that it will manage. For example, Panorama with version 6.0 cannot manage firewalls running PAN-OS 7.0. For versions within the same feature release, although Panorama can manage firewalls running a later version of PAN-OS, Palo Alto Networks recommends that Panorama run the same version or a later version. For example, if Panorama runs 6.0.3, it is recommended that all managed firewalls run PAN-OS 6.0.3 or earlier versions.
- Plan to use the same URL filtering database (BrightCloud or PAN-DB) across all managed firewalls. If some firewalls are using the BrightCloud database and others are using PAN-DB, Panorama can only manage security rules for one or the other URL filtering database. URL filtering rules for the other database must be managed locally on the firewalls that use that database.
- Plan to use Panorama in a high availability configuration; set it up as an active/passive high availability pair. See Panorama High Availability.
- Estimate the log storage capacity your network needs to meet security and compliance requirements. Consider such factors as the network topology, number of firewalls sending logs, type of log traffic (for example, URL Filtering and Threat logs versus Traffic logs), the rate at which firewalls generate logs, and the number of days for which you want to store logs on Panorama. For details, see Determine Panorama Log Storage Requirements.
- For meaningful reports on network activity, plan a logging solution:
  - Do you need to forward logs to a syslog server, in addition to Panorama?
  - If you need a long-term storage solution, do you have a Security Information and Event Management (SIEM) solution, such as Splunk or ArcSight, to which you need to forward logs?
  - Do you need redundancy in logging? With Panorama virtual appliances in HA, each peer can log to its virtual disk. The managed devices can send logs to both peers in the HA pair. This option provides redundancy in logging and is best suited to support up to 2TB of log storage capacity. If you use Dedicated Log Collectors (M-Series appliances in Log Collector mode), you can enable redundancy to ensure that no logs are lost if any one Log Collector in the Collector Group becomes unavailable. Each log will have two copies and each copy will reside on a different Log Collector.
  - Will you log to a Network File System (NFS)? Only the Panorama virtual appliance supports NFS. Consider using NFS if Panorama requires more than 2TB of log storage capacity and but doesn't manage Dedicated Log Collectors. If using NFS, note that the managed devices can send logs only to the primary peer in the HA pair, and only the active-primary Panorama is mounted to the NFS and can write to it.
  - If your logging solution includes M-Series appliances, by default they use the management (MGT) interface for configuration, log collection, and Collector Group communication. However, it is a best practice to use the Eth1 or Eth2 interfaces for log collection and Collector Group communication to improve security, control traffic prioritization, performance, and scalability. Determine whether your solution would benefit from using separate interfaces for these functions. For details, see Set Up the M-Series Appliance.

© Palo Alto Networks, Inc.

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

 © Palo Alto Networks, Inc.

Deploy Panorama: Task Overview

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

Deploy Panorama: Task Overview	
Step 1	(M-Series appliance only) Rack mount the appliance.
Step 2	Perform initial configuration to enable network access to Panorama. See Set Up the Panorama Virtual Appliance or Set Up the M-Series Appliance.
Step 3	Register Panorama and Install Licenses.
Step 4	Install Content and Software Updates for Panorama.
Step 5	(Optional/recommended) Set up Panorama in a high availability configuration. See Panorama High Availability.
Step 6	Add a Firewall as a Managed Device.
Step 7	Add a Device Group or Create a Device Group Hierarchy, Add a Template, and (if applicable) Configure a Template Stack.
Step 8	(Optional) Configure log forwarding to Panorama and/or to external services. See Manage Log Collection.
Step 9	Monitor Network Activity using the visibility and reporting tools on Panorama.

 © Palo Alto Networks, Inc.

Set Up Panorama

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

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

- Determine Panorama Log Storage Requirements
- Set Up the Panorama Virtual Appliance
- Set Up the M-Series Appliance
- Register Panorama and Install Licenses
- Install Content and Software Updates for Panorama
- Transition to a Different Panorama Platform
- Access and Navigate Panorama Management Interfaces
- Set Up Administrative Access to Panorama

Determine Panorama Log Storage Requirements

When you Plan Your Deployment, estimate how much log storage capacity Panorama requires to determine which Panorama Platforms to deploy, whether to expand the storage on those platforms beyond their default capacities, whether to deploy Dedicated Log Collectors, and whether to Configure Log Forwarding from Panorama to External Destinations. When Panorama reaches the maximum capacity, it automatically deletes older logs to create space for new ones. Therefore, to ensure that log retention meets your needs, you should configure any additional storage during the Panorama setup stage. To expand log storage capacity during or after setup, see Expand Log Storage Capacity on the Panorama Virtual Appliance or Increase Storage on the M-Series Appliance.

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

Determine Panorama Log Storage Requirements		
Step 1	Determine the log retention requirements of your organization.	Factors that affect log retention requirements include: • IT policy of your organization • Log redundancy—If you enable log redundancy when you Configure a Collector Group, each log will have two copies, which doubles your required log storage capacity. • Regulatory requirements, such as those specified by the Payment Card Industry Data Security Standard (PCI DSS), Sarbanes-Oxley Act, and Health Insurance Portability and Accountability Act (HIPAA) If your organization requires the removal of logs after a certain period, you can set the expiration period for each log type. You can also set a storage quota for each log type as a percentage of the total space if you need to prioritize log retention by type. For details, see Manage Storage Quotas and Expiration Periods for Logs and Reports.
Step 2	Determine the average daily logging rates. Do this multiple times each day at peak and non-peak times to estimate the average. The more often you sample the rates, the more accurate your estimate.	1. Display the current log generation rate in logs per second: • If Panorama is not yet collecting logs, access the CLI of each firewall, run the following command, and calculate the total rates for all the firewalls. This command displays the number of logs received in the last second. `> debug log-receiver statistics` • If Panorama is already collecting logs, run the following command at the CLI of each platform that receives logs (Panorama management server or Dedicated Log Collector) and calculate the total rates. This command gives the average logging rate for the last five minutes. > **debug log-collector log-collection-stats show incoming-logs** You can also use an SNMP manager to determine the logging rates of M-Series appliances by monitoring the panLcLogRate object (OID 1.3.6.1.4.1.25461.2.3.30.1.1). 2. Calculate the average of the sampled rates. 3. Calculate the daily logging rate by multiplying the average logs-per-second by 86,400.

 © Palo Alto Networks, Inc.

Determine Panorama Log Storage Requirements (Continued)	
Step 3 Estimate the required storage capacity.  This formula provides only an estimate; the exact amount of required storage will differ from the formula result.	Use the formula: <required_storage_duration> x <average_log_size> x <average_logging_rate> The average log size varies considerably by log type. However, you can use 360 bytes as an approximate average log size. For example, if Panorama must store logs for 30 days and the average total logging rate for all firewalls is 21,254,400 logs per day, then the required log storage capacity is: 30 x 360 x 21,254,400 = 229,547,520,000 bytes (approximately 230GB).

Set Up the Panorama Virtual Appliance

The Panorama virtual appliance consolidates the Panorama management and logging functions into a single virtual machine. To enable the use of an existing VMware virtual infrastructure to centrally administer and monitor Palo Alto Networks firewalls, install Panorama on an ESXi server or in vCloud Air.

- ▲ Setup Prerequisites for the Panorama Virtual Appliance
- ▲ Install the Panorama Virtual Appliance
- ▲ Perform Initial Configuration of the Panorama Virtual Appliance
- ▲ Expand Log Storage Capacity on the Panorama Virtual Appliance
- ▲ Increase CPUs and Memory on the Panorama Virtual Appliance
- ▲ Complete the Panorama Virtual Appliance Setup



You can't use the Panorama virtual appliance as a Dedicated Log Collector. You must Set Up the M-Series Appliance in Log Collector mode to have dedicated log collection capabilities. However, you can use the Panorama virtual appliance to manage Dedicated Log Collectors.

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

Setup Prerequisites for the Panorama Virtual Appliance

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

- ❒ Use your browser to access the Palo Alto Networks support site and Register Panorama You will need the Panorama serial number that you received in the order fulfillment email. After registering Panorama, you can access the Panorama software downloads page.
- ❒ If you will install Panorama on a VMware ESXi server, verify that the server meets the following minimum system requirements. These requirements apply to Panorama 5.1 and later releases.

Minimum System Requirements for Panorama on a VMware ESXi Server

- 64-bit kernel-based VMware ESXi 5.1, 5.5, or 6.0. The minimum supported version of the virtual hardware family type (also known as the VMware virtual hardware version) on the ESXi server is vmx-09.
- A client computer with one of the following: VMware vSphere Client or VMware Infrastructure Client that is compatible with your ESXi server.
- Use the following guidelines for allocating CPU and memory:
  - 1–10 managed firewalls: 4 cores and 4GB.
  - 11–50 managed firewalls: 8 cores and 8GB.
  - More than 50 managed firewalls: 8 cores and 16GB.
- 40GB disk space.

  Panorama allocates approximately 11GB for log storage. Increasing the disk space doesn't increase the log storage capacity. To Expand Log Storage Capacity on the Panorama Virtual Appliance, you must add another virtual disk or set up access to a Network File System (NFS) datastore.

 © Palo Alto Networks, Inc.

Install the Panorama Virtual Appliance

- Install Panorama on an ESXi Server
- Install Panorama in vCloud Air

Install Panorama on an ESXi Server

Use these instructions to install a new Panorama virtual appliance on a VMware ESXi server. If you are upgrading your existing Panorama virtual appliance, skip to Install Content and Software Updates for Panorama.

Install Panorama on an ESXi Server	
Step 1 Download the Panorama base image Open Virtual Appliance (OVA) file.	1. Use your browser to access the Palo Alto Networks software downloads site. 2. In the Panorama Base Images section, Download column, click the link for the desired release to download the OVA file.

Install Panorama on an ESXi Server (Continued)	
Step 2 Install Panorama.	1. Launch the VMware vSphere Client and connect to the VMware server. 2. Select **File > Deploy OVF Template**. 3. **Browse** to select the Panorama OVA file and click **Next**. 4. Confirm that the product name and description match the downloaded version, and click **Next**. 5. Enter a descriptive name for the Panorama virtual appliance, and click **Next**. 6. Select a **Datastore Location** on which to install the Panorama image, and click **Next**. Regardless of the disk size, Panorama uses approximately 11GB for log storage. Increasing the disk size doesn't increase the log storage capacity. To Expand Log Storage Capacity on the Panorama Virtual Appliance, you must add another virtual disk or set up access to a Network File System (NFS) datastore. 7. Select **Thick Provision Lazy Zeroed** as the disk format, and click **Next**. 8. Specify which networks in the inventory to use for the Panorama virtual appliance. 9. Confirm the selected options and click **Finish** to start the installation process. 10. When the installation completes, right-click the Panorama virtual appliance and **Edit Settings** as follows: a. Select the **Hardware** tab and allocate **Memory** based on how many firewalls Panorama will manage: – 1-10 firewalls: 4GB – 11-50 firewalls: 8GB – More than firewalls: 16GB b. Set the **SCSI Controller** to **LSI Logic Parallel**. c. Select the **Options** tab, select **General Options**, set the **Guest Operating System** to **Linux**, and set the **Version** to **Other Linux (64-bit)**. d. Click **OK** to save your changes.
Step 3 Power on the Panorama virtual appliance.	In the vSphere Client, right-click the Panorama virtual appliance and select **Power > Power On**. When the Panorama virtual appliance boots, the installation process is complete. You are then ready to Perform Initial Configuration of the Panorama Virtual Appliance.

Install Panorama in vCloud Air

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

 © Palo Alto Networks, Inc.

Install Panorama in vCloud Air	
Step 1 Download the Panorama base image Open Virtual Appliance (OVA) file.	1. Go to the Palo Alto Networks software downloads site. 2. In the Panorama Base Images section, Download column, click the link for the desired release to download the OVA file.

Step 2 Import the Panorama image to the vCloud Air catalog.	For details on these steps, refer to the OVF Tool User's Guide.
	1. Install the OVF Tool on your client system.
	2. Access the client system CLI.
	3. Navigate to the OVF Tool directory (for example, C:\Program Files\VMware\VMware OVF Tool).
	4. Convert the OVA file to an OVF package:
	`ovftool.exe <OVA-file-pathname> <OVF-file-pathname>`
	5. Use a browser to access the vCloud Air web console, select your **Virtual Private Cloud OnDemand** location, and record the browser URL. You will use the URL information to complete the next step. The URL format is: https://<virtual-cloud-location>.vchs.vmware.com/compute/cloud/org/<vCloud-account-number>/#/catalogVAppTemplateList?catalog=<catalog-ID>.
	6. Import the OVF package, using the information from the vCloud Air URL to complete the <virtual-cloud-location>, <vCloud-account-number>, and <catalog-ID> variables. The other variables are your vCloud Air username and domain <user>@<domain>, a virtual data center <datacenter>, and a vCloud Air template <template>.
	`ovftool.exe -st="OVF" "<OVF-file-pathname>" "vcloud://<user>@<domain>:password@<virtual-cloud-location>.vchs.vmware.com?vdc=<datacenter>&org=<vCloud-account-number>&vappTemplate=<template>.ovf&catalog=default-catalog"`

Step 3 Install Panorama.	1. Access the vCloud Air web console and select your **Virtual Private Cloud OnDemand** region. 2. Create a Panorama virtual machine. For the steps, refer to Add a Virtual Machine from a Template in the vCloud Air Documentation Center. Configure the **CPU**, **Memory** and **Storage** as follows: • Set the **CPU** and **Memory** based on how many firewalls Panorama will manage: – 1-10 firewalls—4 CPUs and 4GB memory – 11-50 firewalls—8 CPUs and 8GB memory – More than 50 firewalls—8 CPUs and 16GB memory • Set the **Storage** to at least 40GB. For better logging and reporting performance, select the **SSD-Accelerated** option. Regardless of the disk size, Panorama uses approximately 11GB for log storage. Increasing the disk size doesn't increase the log storage capacity. To increase the log storage capacity, Add a Virtual Disk to Panorama in vCloud Air.

Install Panorama in vCloud Air		
Step 4	Create vCloud Air NAT rules on the gateway to allow inbound and outbound traffic for the Panorama virtual appliance.	Refer to Add a NAT Rule in the vCloud Air Documentation Center for the detailed instructions: 1. Add a NAT rule that allows Panorama to receive traffic from the firewalls and allows administrators to access Panorama. 2. Add a NAT rule that allows Panorama to retrieve updates from the Palo Alto Networks update server and to access the firewalls.
Step 5	Create a vCloud Air firewall rule to allow inbound traffic on the Panorama virtual appliance. Outbound traffic is allowed by default.	Refer to Add a Firewall Rule in the vCloud Air Documentation Center for the detailed instructions.
Step 6	Power on the Panorama virtual appliance if it isn't already on.	In the vCloud Air web console, select the **Virtual Machines** tab, select the Panorama virtual machine, and click **Power On**. You are now ready to Perform Initial Configuration of the Panorama Virtual Appliance.

Perform Initial Configuration of the Panorama Virtual Appliance

Depending on your platform, use the VMware vSphere Client or vCloud Air web console to set up network access to the Panorama virtual appliance. For unified reporting, consider using Greenwich Mean Time (GMT) or Coordinated Universal Time (UTC) as the uniform time zone across Panorama and all the managed devices.

Perform Initial Configuration of the Panorama Virtual Appliance		
Step 1	Gather the required information from your network administrator.	❑ IP address for the management (MGT) interface ❑ Netmask ❑ Default gateway ❑ DNS server IP address
Step 2	Access the console of the Panorama virtual appliance.	1. Access the console. On an ESXi server: a. Launch the VMware vSphere Client. b. Select the **Console** tab for the Panorama virtual appliance and press enter to access the login screen. On vCloud Air: a. Access the vCloud Air web console and select your **Virtual Private Cloud OnDemand** region. b. Select the **Virtual Machines** tab, right-click the Panorama virtual machine, and select **Open In Console**. 2. Enter your username and password to log (default is admin for both). 3. Enter `configure` to switch to Configuration mode.

 © Palo Alto Networks, Inc.

Perform Initial Configuration of the Panorama Virtual Appliance (Continued)	

Step 3 Configure the network access settings for the MGT interface.
Panorama uses the MGT interface for management traffic, high availability synchronization, log collection, and communication within Collector Groups.
	1. Enter the following command, where <Panorama-IP> is the IP address you want to assign to the Panorama management interface, <netmask> is the subnet mask, <gateway-IP> is the IP address of the network gateway, and <DNS-IP> is the IP address of the DNS server:
	set deviceconfig system ip-address <Panorama-IP> **netmask** <netmask> **default-gateway** <gateway-IP> **dns-setting servers primary** <DNS-IP>
	2. Enter **commit** to commit your changes.
	3. Enter **exit** to exit Configuration mode.
	4. Use the ping utility to verify network access to external services required for firewall management, such as the default gateway, DNS server, and the Palo Alto Networks Update Server, as shown in the following example:
	`admin@Panorama-Corp> ping host updates.paloaltonetworks.com`
	`PING updates.paloaltonetworks.com (67.192.236.252) 56(84) bytes of data.`
	`64 bytes from 67.192.236.252: icmp_seq=1 ttl=243 time=40.5 ms`
	`64 bytes from 67.192.236.252: icmp_seq=1 ttl=243 time=53.6 ms`
	`64 bytes from 67.192.236.252: icmp_seq=1 ttl=243 time=79.5 ms`
	After verifying connectivity, press Ctrl+C to stop the pings.

Step 4 Configure the general settings.	1. Using a secure connection (HTTPS) from a web browser, log in to the Panorama web interface using the IP address and password you assigned to the management interface (https://<IP address>). 2. Select **Panorama > Setup > Management** and edit the General Settings. 3. Enter a **Hostname** for the server and enter the network **Domain** name. The domain name is just a label; Panorama doesn't use it to join the domain. 4. Align the clock on Panorama and the managed firewalls to use the same **Time Zone**, for example GMT or UTC. Timestamps are recorded when Panorama receives the logs and the managed firewalls generate the logs. Aligning the time zones on Panorama and the firewalls ensures that the timestamps are synchronized and the process of querying logs and generating reports on Panorama is harmonious. 5. Enter the **Latitude** and **Longitude** to enable accurate placement of the Panorama management server on the world map. 6. Enter the **Serial Number** you received in the order fulfillment email. 7. Click **OK** to save your changes.
Step 5 Change the default administrator password. To ensure that the management interface remains secure, configure the Minimum Password Complexity (**Panorama > Setup > Management**).	1. Click the **admin** link on the left side of the web interface footer. 2. Enter the **Old Password** and the **New Password** in the appropriate fields and record the new password in a safe location. 3. Click **OK**.

Perform Initial Configuration of the Panorama Virtual Appliance (Continued)		
Step 6	(Optional) Modify the management interface settings.	1. Select **Panorama > Setup > Management** and edit the Management Interface Settings. 2. Select which management services to allow on the interface (for example, Secure Shell (**SSH**) access) As a best practice, don't select **Telnet** or **HTTP**. These services use plaintext and are less secure than the other services. 3. Click **OK**.
Step 7	Save your configuration changes.	Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. If necessary, you can Expand Log Storage Capacity on the Panorama Virtual Appliance. Otherwise, you are ready to Complete the Panorama Virtual Appliance Setup.

Expand Log Storage Capacity on the Panorama Virtual Appliance

By default, the Panorama virtual appliance has a single disk partition for all data in which, regardless of the total disk size, approximately 11GB is allocated for log storage. Increasing the disk size doesn't increase the log storage capacity. If you need up to 2TB of log storage, you can add a virtual disk to Panorama installed on a VMware ESXi server or in VMware vCloud Air. If you need more than 2TB, you can mount Panorama to an NFS datastore but only on the ESXi server, not in vCloud Air. Before expanding log storage capacity, Determine Panorama Log Storage Requirements.

For additional log storage, you can also forward firewall logs to Dedicated Log Collectors (see Configure a Collector Group) or Configure Log Forwarding from Panorama to External Destinations.

- Add a Virtual Disk to Panorama on an ESXi Server
- Add a Virtual Disk to Panorama in vCloud Air
- Mount the Panorama ESXi Server to an NFS Datastore

Add a Virtual Disk to Panorama on an ESXi Server

To expand log storage capacity beyond the approximately 11GB internal storage allocated by default on the Panorama virtual appliance, you can add another virtual disk of up to 2TB.

If Panorama loses connectivity to the new virtual disk, Panorama might lose logs during the failure interval.

To allow for redundancy, use the virtual disk in a RAID configuration. RAID10 provides the best write performance for applications with high logging characteristics.

If necessary, you can Replace the Virtual Disk on a Panorama ESXi Server.

 © Palo Alto Networks, Inc.

Add a Virtual Disk to Panorama on an ESXi Server	
Step 1	Access the VMware vSphere Client and select the **Virtual Machines** tab.
Step 2	Right-click the Panorama virtual appliance and select **Power > Power Off**.
Step 3	Right-click the Panorama virtual appliance and select **Edit Settings**.
Step 4	Click **Add** in the **Hardware** tab to launch the Add Hardware wizard.
Step 5	Select **Hard Disk** for the hardware type and click **Next**.
Step 6	Select **Create a new virtual disk** and click **Next**.
Step 7	Set the **Disk Size** to 2TB.
Step 8	Select the **Thick Provision Lazy Zeroed** disk format.
Step 9	Set the **Location** to **Store with the virtual machine** (the datastore doesn't have to reside on the ESXi server) and click **Next**.
Step 10	Select a SCSI **Virtual Device Node** (you can use the default selection) and click **Next**. The selected node must be in SCSI format; Panorama will fail to boot if you select another format.
Step 11	Verify that the settings are correct and click **Finish** to exit the wizard. The new disk appears in the list of devices for the virtual appliance.
Step 12	Right-click the Panorama virtual appliance and select **Power > Power On**. The virtual disk initializes for first-time use. The size of the new disk determines how long initialization takes. After initialization, Panorama moves all existing logs on the internal storage to the new disk and writes all new entries to it.
Step 13	Log in to Panorama, select **Panorama > Setup > Management** and, in the Logging and Reporting Settings section, verify that the **Log Storage** capacity accurately displays the new disk capacity.

Add a Virtual Disk to Panorama in vCloud Air

To expand log storage capacity beyond the approximately 11GB internal storage allocated by default on the Panorama virtual appliance in vCloud Air, you can add another virtual disk of up to 2TB.

If Panorama loses connectivity to the new virtual disk, Panorama might lose logs during the failure interval.
If necessary, you can Replace the Virtual Disk on vCloud Air.

Add a Virtual Disk to Panorama on an ESXi Server	
Step 1	Access the vCloud Air web console and select your **Virtual Private Cloud OnDemand** region.
Step 2	Select the Panorama virtual appliance in the **Virtual Machines** tab.
Step 3	Select **Actions > Edit Resources** and **Add another disk**.
Step 4	Set the **Storage** to 2TB and specify the storage tier (**Standard** or **SSD-Accelerated**).
Step 5	**Save** your changes.

Add a Virtual Disk to Panorama on an ESXi Server (Continued)	
Step 6	Log in to Panorama, select **Panorama > Setup > Management** and, in the Logging and Reporting Settings section, verify that the **Log Storage** capacity accurately displays the new disk capacity.

Mount the Panorama ESXi Server to an NFS Datastore

When the Panorama virtual appliance runs on an ESXi server, mounting to a Network File Share (NFS) datastore provides the ability to write logs to a centralized location and offers the flexibility to expand the log storage capacity beyond 2TB. Before setting up an NFS datastore in a Panorama high availability configuration, see Logging Considerations in Panorama HA.

Mount the Panorama ESXi Server to an NFS Datastore	
Step 1	Select **Panorama > Setup > Operations** and, in the Miscellaneous section, click **Storage Partition Setup**.
Step 2	Set the **Storage Partition** type to **NFS V3**.
Step 3	Enter the IP address of the NFS **Server**.
Step 4	Enter the **Log Directory** path for storing the log files. For example, export/panorama.
Step 5	For the **Protocol**, select **TCP** or **UDP**, and enter the **Port** for accessing the NFS server. To use NFS over TCP, the NFS server must support it. Common NFS ports are UDP/TCP 111 for RPC and UDP/TCP 2049 for NFS.
Step 6	For optimal NFS performance, in the **Read Size** and **Write Size** fields, specify the maximum size of the chunks of data that the client and server pass back and forth to each other. Defining a read/write size optimizes the data volume and speed in transferring data between Panorama and the NFS datastore.
Step 7	(Optional) Select **Copy On Setup** to copy the existing logs stored on Panorama to the NFS volume. If Panorama has a lot of logs, this option might initiate the transfer of a large volume of data.
Step 8	Click **Test Logging Partition** to verify that Panorama can access the NFS **Server** and **Log Directory**.
Step 9	Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again. Until you reboot, the Panorama virtual appliance writes logs to the local storage disk.
Step 10	Select **Panorama > Setup > Operations** and select **Reboot Panorama** in the Device Operations section. After rebooting, Panorama starts writing logs to the NFS datastore.

Increase CPUs and Memory on the Panorama Virtual Appliance

When you Perform Initial Configuration of the Panorama Virtual Appliance, you specify the memory and number of CPUs based on how many firewalls Panorama will manage:

- 1–10 firewalls—4 CPUs and 4GB memory
- 11–50 firewalls—8 CPUs and 8GB memory
- 51–1,000 firewalls—8 CPUs and 16GB memory

If you increase the number of firewalls that Panorama manages, you must also increase the memory and CPUs.

 © Palo Alto Networks, Inc.

▲ Increase CPUs and Memory for Panorama on an ESXi Server
▲ Increase CPUs and Memory for Panorama on vCloud Air

Increase CPUs and Memory for Panorama on an ESXi Server

Increase CPUs and Memory for Panorama on an ESXi Server	
Step 1	Access the VMware vSphere Client and select **Virtual Machines**.
Step 2	Right-click the Panorama virtual appliance and select **Power > Power Off**.
Step 3	Right-click the Panorama virtual appliance and select **Edit Settings**.
Step 4	Select **Memory** and enter the new **Memory Size** based on the number of firewalls that Panorama manages: • 1-10 firewalls—4GB memory • 11-50 firewalls—8GB memory • 51-1,000 firewalls—16GB memory
Step 5	Select **CPUs** and specify the number of CPUs (the **Number of virtual sockets** multiplied by the **Number of cores per socket**) based on the number of firewalls that Panorama manages: • 1-10 firewalls—4 CPUs • 11-50 firewalls—8 CPUs • 51-1,000 firewalls—8 CPUs
Step 6	Click **OK** to save your changes.
Step 7	Right-click the Panorama virtual appliance and select **Power > Power On**.

Increase CPUs and Memory for Panorama on vCloud Air

Increase CPUs and Memory for Panorama on vCloud Air	
Step 1	Access the vCloud Air web console and select your **Virtual Private Cloud OnDemand** region.
Step 2	In the **Virtual Machines** tab, select the Panorama virtual machine and **Power Off**.
Step 3	Select **Actions > Edit Resources**.
Step 4	Set the **CPU** and **Memory** based on the number of firewalls that Panorama manages: • 1–10 firewalls—4 CPUs and 4GB memory • 11–50 firewalls—8 CPUs and 8GB memory • 51–1,000 firewalls—8 CPUs and 16GB memory
Step 5	**Save** your changes.
Step 6	Select the Panorama virtual machine and **Power On**.

Complete the Panorama Virtual Appliance Setup

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

- Activate a Panorama Support License
- Activate/Retrieve a Device Management License on the Panorama Virtual Appliance
- Install Content and Software Updates for Panorama
- Access and Navigate Panorama Management Interfaces
- Set Up Administrative Access to Panorama
- Manage Firewalls

 © Palo Alto Networks, Inc.

Set Up the M-Series Appliance

The M-Series appliance is a high performance hardware platform that you can deploy in Panorama mode or Log Collector mode. When you Perform Initial Configuration of the M-Series Appliance, you can configure the following interfaces. The M-100 and M-500 Appliance Hardware Reference Guides explain where to attach cables for these interfaces.



The M-Series appliance does not support Link Aggregation Control Protocol (LACP) for aggregating these interfaces.

The Eth1 and Eth2 interfaces are available only if the M-Series appliance runs Panorama 6.1 or a later release and the managed firewalls run PAN-OS 5.0 or a later release.

Interface	Description
Management (MGT)	You must use the MGT interface for management and configuration of the firewalls, Log Collectors, and Panorama. By default, Panorama also uses MGT for log collection and communication within Collector Groups, though you can reassign these functions to the Eth1 and Eth2 interfaces.
Eth1 Eth2	You can configure the M-Series appliance to use Eth1 or Eth2 for log collection and Collector Group communication. Each interface can support either or both of these functions—for example, you can configure Eth1 for both log collection and Collector Group communication. You cannot, however, assign multiple interfaces for the same function—for example, you cannot configure both Eth1 and Eth2 for log collection.  To reduce the traffic load on MGT and improve security for management traffic, use Eth1 and/or Eth2 for log collection and Collector Group communication.
Eth3	Reserved for future use.

Use the following workflows for setting up an M-Series appliance:

M-Series Appliance in Panorama Mode	M-Series Appliance in Log Collector Mode
Step 1 Rack mount the M-Series appliance. Refer to the M-100 or M-500 Appliance Hardware Reference Guide for instructions.	Step 1 Rack mount the M-Series appliance. Refer to the M-100 or M-500 Appliance Hardware Reference Guide for instructions.
Step 2 Perform Initial Configuration of the M-Series Appliance	Step 2 Perform Initial Configuration of the M-Series Appliance
Step 3 Register Panorama and Install Licenses	Step 3 Register Panorama and Install Licenses
Step 4 Install Content and Software Updates for Panorama	Step 4 Install Content and Software Updates for Panorama
Step 5 Configure each logging disk pair. This task is required to make the RAID disks available for logging. Optionally, you can add disks to Increase Storage on the M-Series Appliance.	Step 5 Configure each logging disk pair. This task is required to make the RAID disks available for logging. Optionally, you can add disks to Increase Storage on the M-Series Appliance.
Step 6 Set Up Administrative Access to Panorama	Step 6 Set up the M-Series Appliance as a Log Collector
Step 7 Manage Firewalls	Step 7 Manage Log Collection
Step 8 Manage Log Collection	

Perform Initial Configuration of the M-Series Appliance

By default, Panorama has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other configuration tasks. You must perform these initial configuration tasks either from the Management (MGT) interface or using a direct serial port connection to the console port on the M-Series appliance.

Perform Initial Configuration of the M-Series Appliance		
Step 1	Gather the required interface and server information from your network administrator.	❑ Gather the IP address, netmask (for IPv4) or prefix length (for IPv6), and default gateway for each interface that you plan to configure (MGT, Eth1, and/or Eth2). Only the MGT interface is mandatory.  If you plan to use the appliance as a Panorama management server, Palo Alto Networks recommends that you configure the Eth1 and/or Eth2 interfaces for log collection and Collector Group communication (see Set Up the M-Series Appliance). You can improve the security of management traffic by defining a separate subnet for the MGT interface that is more private than the Eth1 and Eth2 subnets. ❑ Gather the IP addresses of the DNS servers.
Step 2	Access the M-Series appliance from your computer.	1. Connect to the M-Series appliance in one of the following ways: • Attach a serial cable from a computer to the Console port on the M-Series appliance and connect using terminal emulation software (9600-8-N-1). • Attach an RJ-45 Ethernet cable from a computer to the MGT port on the M-Series appliance. From a browser, go to **https://192.168.1.1.** Enabling access to this URL might require changing the IP address on the computer to an address in the 192.168.1.0 network (for example, 192.168.1.2). 2. When prompted, log in to the appliance using the default username and password (admin/admin). The appliance starts initializing.

 © Palo Alto Networks, Inc.

Perform Initial Configuration of the M-Series Appliance (Continued)		
Step 3	Configure the network access settings for each interface that Panorama will use for configuration, log collection, and Collector Group communication.	1. Select **Panorama > Setup > Management**. 2. Edit the Interface Settings of each interface that Panorama will use. Only the MGT interface is required. The Eth1 and Eth2 interfaces are optional and apply only if you plan to use the M-Series appliance as a Panorama management server. a. Complete one or both of the following field sets, depending on the IP protocols of your network: – IPv4—**IP Address**, **Netmask**, and **Default Gateway** – IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway** b. (Optional) Select the check boxes for the management services to allow on the interface. As a best practice, clear the **Telnet** and **HTTP** check boxes for the MGT interface: these services use plaintext and so are less secure than others. c. Click **OK** to save your changes.
Step 4	Configure the hostname, time zone, and general settings.	1. Select **Panorama > Setup > Management** and edit the General Settings. 2. Align the clock on Panorama and the managed firewalls to use the same **Time Zone**, for example GMT or UTC. The firewall records timestamps when it generate logs and Panorama records timestamps upon receiving the logs. Aligning the time zones ensures that the timestamps are synchronized and that the process of querying logs and generating reports on Panorama is harmonious. 3. Enter a **Hostname** for the server. Panorama uses this as the display name/label for the appliance. For example, this is the name that appears at the CLI prompt. It also appears in the Collector Name field if you add the appliance as a managed collector on the **Panorama > Managed Collectors** page. 4. (Optional) Enter the **Latitude** and **Longitude** to enable accurate placement of the M-Series appliance on the world map. The **App Scope > Traffic Maps** and **App Scope > Threat Maps** use these values. 5. Click **OK** to save your entries.
Step 5	Configure the DNS and update servers.	1. Select **Panorama > Setup > Services** and edit the settings. 2. Enter the IP address of the **Primary DNS Server** and (optionally) of the **Secondary DNS Server**. 3. Enter the URL or static address of the **Update Server** (default updates.paloaltonetworks.com). Select **Verify Update Server Identity** if you want Panorama to verify that the Update Server from which it downloads software or content packages has an SSL certificate that a trusted authority signed. This option adds an additional level of security for communication between the Panorama management server and Update Server. 4. Click **OK** to save your entries.

Perform Initial Configuration of the M-Series Appliance (Continued)	
Step 6 Change the default admin password.	1. Click the **admin** link in the lower left of the web interface. 2. Enter the **Old Password**, **New Password**, and **Confirm New Password**, and then click **OK**. Store the new password in a safe location.  To ensure that the MGT interface remains secure, configure Minimum Password Complexity settings (select **Panorama > Setup > Management**) and specify the interval at which administrators must change their passwords.
Step 7 Commit your configuration changes.	Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again.  If you plan to use the M-Series appliance as a Panorama management server and you configured the Eth1 or Eth2 interfaces, you must assign them to the **Device Log Collection** or **Collector Group Communication** functions when you Configure a Managed Collector. To make the interfaces operational, you must then Configure a Collector Group for the managed collector and perform a Collector Group commit.

 © Palo Alto Networks, Inc.

Perform Initial Configuration of the M-Series Appliance (Continued)		

Step 8	Verify network access to external services required for Panorama management, such as the Palo Alto Networks Update Server.	1. Connect to the M-Series appliance in one of the following ways:
		• Attach a serial cable from your computer to the Console port on the M-Series appliance. Then use a terminal emulation software (9600-8-N-1) to connect.
		• Use terminal emulation software such as PuTTY to open an SSH session to the IP address that you specified for the MGT interface of the M-Series appliance during initial configuration.
		2. Log in to the CLI when prompted. Use the default admin account and the password that you specified during initial configuration.
		3. Use the ping utility to verify network connectivity to the Palo Alto Networks Update Server as shown in the following example. Verify that DNS resolution occurs and the response includes the IP address for the Update Server (10.101.16.13, in this example); the Update Server does not respond to a ping request.
		`> ping host updates.paloaltonetworks.com`
		`PING updates.paloaltonetworks.com (10.101.16.13) 56(84) bytes of data.`
		After verifying DNS resolution, press Ctrl+C to stop the ping request.
		4. Use the following CLI command to retrieve information on the support entitlement for Panorama from the Update Server:
		`> request support check`
		If you have connectivity, the Update Server responds with the support status for Panorama. Because Panorama is not registered, the Update Server returns the following message:
		`Contact Us https://www.paloaltonetworks.com/company/contact-us.html Support Home https://www.paloaltonetworks.com/support/tabs/overview.html Device not found on this update server`

Step 9	Next steps...	1. Register Panorama and Install Licenses. 2. Install Content and Software Updates for Panorama.

Set up the M-Series Appliance as a Log Collector

If you want a dedicated appliance for log collection, configure an M-100 or M-500 appliance in Log Collector mode. To do this, you first perform the initial configuration of the appliance in Panorama mode, which includes licensing, installing software and content updates, and configuring the management (MGT) interface. You then switch the M-100 or M-500 appliance to Log Collector mode and complete the Log Collector configuration. Additionally, if you want to use dedicated interfaces (recommended) instead of the

MGT interface for log collection and Collector Group communication, you must first configure the interfaces for the Panorama management server, then configure them for the Log Collector, and then perform a Panorama commit followed by a Collector Group commit.

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



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

Set up the M-100 or M-500 Appliance as a Log Collector		
Step 1	Set up the Panorama management server that will manage the Log Collector if you have not already done so.	Perform one of the following tasks: • Set Up the Panorama Virtual Appliance • Set Up the M-Series Appliance
Step 2	Record the management IP addresses of the Panorama management server. If you deployed Panorama in a high availability (HA) configuration, you need the IP address of each HA peer.	1. Log in to the web interface of the Panorama management server. 2. Record the **IP Address** of the solitary (non-HA) or active (HA) Panorama by selecting **Panorama > Setup > Management** and checking the Management Interface Settings. 3. For an HA deployment, record the **Peer HA IP Address** of the passive Panorama by selecting **Panorama > High Availability** and checking the Setup section.
Step 3	Set up the M-Series appliance that will serve as a Dedicated Log Collector. If you previously deployed this appliance as a Panorama management server, you can skip this step because the MGT interface is already configured and the licenses and updates are already installed.	The M-Series appliance in Log Collector mode does not have a web interface for configuration tasks, only a CLI. Therefore, before changing the mode on the M-Series appliance, use the web interface in Panorama mode to: 1. Perform Initial Configuration of the M-Series Appliance. 2. Register Panorama and Install Licenses. 3. Install Content and Software Updates for Panorama.
Step 4	Access the CLI of the M-Series appliance.	1. Connect to the M-Series appliance in one of the following ways: • Attach a serial cable from your computer to the Console port on the M-Series appliance. Then use terminal emulation software (9600-8-N-1) to connect. • Use terminal emulation software such as PuTTY to open an SSH session to the IP address that you specified for the MGT interface of the M-Series appliance during initial configuration. 2. Log in to the CLI when prompted. Use the default admin account and the password that you specified during initial configuration.

 © Palo Alto Networks, Inc.

Set up the M-100 or M-500 Appliance as a Log Collector (Continued)		

Step 5	Switch from Panorama mode to Log Collector mode.	1. Switch to Log Collector mode by entering the following command:
		`> request system system-mode logger`
		2. Enter **Y** to confirm the mode change. The M-Series appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the M-Series appliance to see the Panorama login prompt.
		If you see a **CMS Login** prompt, this means the Log Collector has not finished rebooting. Press Enter at the prompt without typing a username or password.
		3. Log back in to the CLI.
		4. Verify that the switch to Log Collector mode succeeded:
		`> show system info | match system-mode`
		If the mode change succeeded, the output displays:
		`system-mode: logger`

Step 6	Configure the logging disks as RAID1 pairs.
	If you previously deployed the appliance as a Panorama management server, you can skip this step because the disk pairs are already configured and available.
	The time required to configure the drives varies from several minutes to a couple of hours, based on the amount of data on the drives.
		1. Determine which disk pairs are present for configuring as RAID pairs on the M-Series appliance:
		`> show system raid detail`
		Perform the remaining steps to configure each disk pair that has `present` disks. This example uses disk pair A1/A2.
		2. To add the first disk in the pair, enter the following command and enter **y** when prompted to confirm the request:
		`> request system raid add A1`
		Wait for the process to finish before adding the next disk in the pair. To monitor the progress of the RAID configuration, re-enter:
		`> show system raid detail`
		After the process finishes for the first disk, the output displays the disk pair status as `Available` but `degraded`.
		3. Add the second disk in the pair:
		`> request system raid add A2`
		4. Verify that the disk setup is complete:
		`> show system raid detail`
		After the process finishes for the second disk, the output displays the disk pair status as `Available` and `clean`:
		`Disk Pair A Available`
		`Status clean`

Step 7	Record the serial number of the Log Collector. You need the serial number to add the Log Collector as a managed collector on the Panorama management server.	1. At the Log Collector CLI, enter the following command to display its serial number. `> show system info | match serial` 2. Record the serial number.

Set up the M-100 or M-500 Appliance as a Log Collector (Continued)		
Step 8	Add the Log Collector as a managed collector to the Panorama management server.	1. Select **Panorama > Managed Collectors** and **Add** a managed collector. 2. In the **General** tab, enter the serial number (**Collector S/N**) you recorded for the Log Collector. 3. In the **Panorama Server IP** field, enter the IP address or FQDN of the solitary (non-HA) or active (HA) Panorama. For HA deployments, enter the IP address or FQDN of the passive Panorama peer in the **Panorama Server IP 2** field. 4. Select **Management** and configure one or both of the following field sets for the MGT interface based on the IP protocols of your network. • IPv4—**IP Address**, **Netmask**, and **Default Gateway** • IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway** 5. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again. This step is required before you can enable logging disks. 6. Verify that the **Panorama > Managed Collectors** page lists the Log Collector you added. The Connected column displays a check mark to indicate that the Log Collector is connected to Panorama. You might have to wait a few minutes before the page displays the updated connection status. At this point, the Configuration Status column displays Out of Sync and the Run Time Status column displays disconnected. The status will change to In Sync and connected after you configure a Collector Group (Step 12).

Step 9	Enable connectivity between the Log Collector and Panorama management server.	Enter the following commands at the Log Collector CLI, where `<IPaddress1>` is for the MGT interface of the solitary (non-HA) or active (HA) Panorama and `<IPaddress2>` is for the MGT interface of the passive (HA) Panorama, if applicable.
		`> configure`
		`# set deviceconfig system panorama-server <IPaddress1> panorama-server-2 <IPaddress2>`
		`# commit`
		`# exit`

Step 10	Enable the logging disks.	1. Select **Panorama > Managed Collectors** and edit the Log Collector. 2. Select **Disks** and **Add** each RAID disk pair. 3. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again.

 © Palo Alto Networks, Inc.

Set up the M-100 or M-500 Appliance as a Log Collector (Continued)	
Step 11 (Optional) Configure the **Eth1** and/or **Eth2** interfaces if the Panorama management server and Log Collector will use them for log collection and Collector Group communication. If you previously deployed the Log Collector as a Panorama management server and configured the **Eth1** and/or **Eth2** interfaces, you must reconfigure those interfaces because switching to Log Collector mode (Step 5) would have deleted all configurations except the management access settings.  Palo Alto Networks recommends using **Eth1** and/or **Eth2** to reduce the traffic load on the MGT interface and to improve security for management traffic.	1. Configure **Eth1** and/or **Eth2** on the Panorama management server if you haven't already: a. Select **Panorama > Setup > Management**. b. Edit the Eth1 Interface Settings and/or Eth2 Interface Settings. For each interface, complete one or both of the following field sets based on the IP protocols of your network: – IPv4—**IP Address**, **Netmask**, and **Default Gateway** – IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway** c. Click **OK** to save your changes. 2. Configure **Eth1** and/or **Eth2** on the Log Collector: a. Select **Panorama > Managed Collectors** and edit the Log Collector. b. Configure the network settings of the **Eth1** and/or **Eth2** interfaces. For each interface, select the corresponding tab and configure one or both of the following field sets based on the IP protocols of your network. – IPv4—**IP Address**, **Netmask**, and **Default Gateway** – IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway** c. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again. This step is required before you can assign the Eth1 and Eth2 interfaces to logging functions. d. Edit the Log Collector again and select the interfaces (MGT, Eth1, or Eth2) that it will use for **Device Log Collection** and **Collector Group Communication** (default is MGT). e. Click **OK** to save your changes.

Set up the M-100 or M-500 Appliance as a Log Collector (Continued)	

Step 12 Assign the Log Collector to a Collector Group.	1. Configure a Collector Group. You must perform a Panorama commit and then a Collector Group commit to synchronize the Log Collector configuration with Panorama and to put the Eth1 and Eth2 interfaces (if you configured them) in an operational state on the Log Collector.
	All the Log Collectors for any particular Collector Group must be the same platform: all M-100 appliances or all M-500 appliances. If you **Enable log redundancy across collectors** in the Collector Group, then each Log Collector also requires the same number of disks.
	2. Select **Panorama > Managed Collectors** to verify that the Log Collector configuration is synchronized with Panorama.
	The Configuration Status column should display In Sync and the Run Time Status column should display connected.
	3. Access the Log Collector CLI and enter the following command to verify that its interfaces are operational:
	`> show interface all`
	The output displays the `state` as `up` for each interface that is operational.
	4. If the Collector Group has multiple Log Collectors, verify they can communicate with each other by running the following command for each interface that the Log Collectors use (MGT, Eth1, and/or Eth2). For the `source` IP address, specify the interface of the Log Collector on which you run the command. For the `host` IP address, specify the matching interface of another Log Collector in the same Collector Group.
	`> ping source <IP-address> host <IP-address>`
	For example, if a Collector Group contains Log Collector A with an MGT interface set to 192.0.2.1 and Log Collector B with an MGT interface set to 192.0.2.2, log in to Log Collector A and enter:
	`> ping source 192.0.2.1 host 192.0.2.2`
	If the Log Collectors can communicate over their MGT interfaces, the output displays:
	`PING 192.0.2.2 (192.0.2.2) from 192.0.2.1 : 56(84) bytes of data.`

Step 13 Next steps...	To enable the Log Collector to receive firewall logs: 1. Configure Log Forwarding to Panorama. 2. Verify Log Forwarding to Panorama.

Increase Storage on the M-Series Appliance

The M-100 appliance ships with two disks for 1TB of storage, and you can add up to three disk pairs (3TB) to reach the maximum of 4TB of storage. The M-500 appliance ships with eight disks for 4TB of storage, and you can add up to four disk pairs (4TB) to reach the maximum of 8TB of storage. In both appliances, the disks have a RAID 1 configuration.

 © Palo Alto Networks, Inc.

Before expanding log storage capacity, Determine Panorama Log Storage Requirements.



If you need more log storage than a single M-Series appliance supports, you can add more M-Series appliances in Log Collector mode (see Configure a Managed Collector) or Configure Log Forwarding from Panorama to External Destinations.

If adding disk pairs to an already deployed M-Series appliance, you don't need to take the appliance offline to expand the storage. When the additional disk pairs become available, the M-Series appliance redistributes the logs among the disk pairs. This log redistribution process happens in the background and doesn't impact uptime or the availability of the M-Series appliance. However, the process does diminish the maximum logging rate. In the **Panorama > Collector Groups** page, the Redistribution State column indicates the completion status of the process as a percentage.

Increase Storage on the M-Series Appliance		
Step 1	Install the new disks in the appropriate drive bays.	Make sure to add the drives sequentially in the next open disk bay slot for the disk pair. For example, add B1/B2 before C1/C2. For information on adding the physical drives, refer to the M-100 or M-500 Hardware Reference Guide.
Step 2	Access the command line interface (CLI) on the M-Series appliance.	You can connect to the M-Series appliance in one of the following ways: • Connect a serial cable from your computer to the Console port and connect to the M-Series appliance using terminal emulation software (9600-8-N-1). • Use a terminal emulation software such as PuTTY to open a Secure Shell (SSH) session to the IP address of the M-Series appliance.
Step 3	When prompted, log in to the appliance.	Use the default admin account and the password assigned.

Increase Storage on the M-Series Appliance (Continued)	

Step 4 Configure each logging disk pair.
  The time required to mirror the data on the drive may vary from several minutes to a couple hours, depending on the amount of data on the drive.
	This example uses the drives in the disk bays B1 and B2.
	1. Enter the following commands and confirm the request when prompted:
	request system raid add B1
	request system raid add B2
	2. To monitor the progress of the RAID configuration, enter the following command:
	show system raid detail
	When the RAID set up is complete, the following response displays:
	`Disk Pair A Available`
	`Status clean`
	`Disk id A1 Present`
	`model : ST91000640NS`
	`size : 953869 MB`
	`status : active sync`
	`Disk id A2 Present`
	`model : ST91000640NS`
	`size : 953869 MB`
	`status : active sync`
	`Disk Pair B Available`
	`Status clean`
	`Disk id B1 Present`
	`model : ST91000640NS`
	`size : 953869 MB`
	`status : active sync`
	`Disk id B2 Present`
	`model : ST91000640NS`
	`size : 953869 MB`
	`status : active sync`

Step 5 Make the disk pair available for logging. To enable the disk pairs for logging, this appliance must have been added as a managed collector on Panorama. If you have not already added it, see Manage Collector Groups.	1. Access the Panorama management server that is managing this Log Collector (if it is a different appliance). 2. On the **Panorama > Managed Collectors** tab, select the Log Collector and follow the instructions in Step 9 in Manage Collector Groups. 3. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

 © Palo Alto Networks, Inc.

Register Panorama and Install Licenses

Before you can begin using Panorama for centralized management, logging, and reporting, you must register, activate, and retrieve the Panorama licenses. Every instance of Panorama requires valid licenses that entitle you to manage firewalls and obtain support. The device management license enforces the maximum number of firewalls that Panorama can manage. This license is based on firewall serial numbers, not on the number of virtual systems on each firewall. The support license enables Panorama software updates and dynamic content updates (for the latest Applications and Threats signatures, as an example). To purchase licenses, contact your Palo Alto Networks Systems Engineer or reseller.

- Register Panorama
- Activate a Panorama Support License
- Activate/Retrieve a Device Management License on the Panorama Virtual Appliance
- Activate/Retrieve a Device Management License on the M-Series Appliance



If you are running an evaluation license for device management on your Panorama virtual appliance and want to apply a Panorama license that you purchased, perform the tasks Register Panorama and Activate/Retrieve a Device Management License on the Panorama Virtual Appliance.

Register Panorama

Register Panorama		
Step 1	Record the Panorama serial number or auth-code and record your Sales Order Number or Customer ID.	For the auth-code, Sales Order Number, or Customer ID, see the order fulfillment email that Palo Alto Networks Customer Service sent when you placed your order for Panorama. For the serial number, the location depends on the platform: • M-Series appliance—Log in to the Panorama web interface and record the **Serial #** value in the **Dashboard** tab, General Information section. • Panorama virtual appliance—See the order fulfillment email.

Register Panorama (Continued)	
Step 2 Register Panorama. The steps depend on whether you already have a login for the Support site.	• If this is the first Palo Alto Networks appliance you are registering and you don't yet have a login: a. Go to the Palo Alto Networks Support site. b. Click **Register** on the right side of the page, enter your **Email Address**, enter the code displayed on the page, and click **Submit**. c. Complete the fields in the **Create Contact Details** section. d. Enter a **Display Name**, **Confirm Email Address**, and **Password/Confirm Password**. e. Enter the Panorama **Device Serial Number** or **Auth Code**. f. Enter your **Sales Order Number** or **Customer ID**. g. Click **Submit**. • If you already have a support account: a. Log in to the Palo Alto Networks Support site. b. Click the **Assets** tab, and click **Register New Device**. c. Enter the Panorama **Device Serial Number**. d. Enter your **City**, **Postal Code**, and **Country**. e. Click **Submit**.

Activate a Panorama Support License

Before activating a Panorama support license on a Panorama M-Series appliance or Panorama virtual appliance, you must Register Panorama.

If the support license expires, Panorama can still manage firewalls and collect logs, but software and content updates will be unavailable. The software and content versions on Panorama must be the same as or later than the versions on the managed firewalls, or else errors will occur. For details, see Panorama, Log Collector, and Firewall Version Compatibility.

Activate a Panorama Support License	
Step 1	Select **Panorama > Support** and click **Activate feature using authorization code**.
Step 2	Enter the **Authorization Code** and click **OK**.
Step 3	Verify that the subscription is activated. 

 © Palo Alto Networks, Inc.

Activate/Retrieve a Device Management License on the Panorama Virtual Appliance

Before activating and retrieving a device management license on the Panorama virtual appliance, you must Register Panorama. If you are running an evaluation license and want to apply a license that you purchased, you must still register and activate/retrieve the purchased license.

Activate/Retrieve a Device Management License on the Panorama Virtual Appliance	
Step 1	Select **Panorama > Setup > Management** and edit the General Settings.
Step 2	Enter the Panorama **Serial Number** (included in the order fulfillment email) and click **OK**.
Step 3	Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. To determine how many firewalls a license enables the Panorama virtual appliance to manage, log in to the Palo Alto Support website (https://support.paloaltonetworks.com), select the **Assets** tab, find the Panorama device, and view the Model Name. For example, a license for the PAN-PRA-25 model can manage 25 devices. This page also displays the Expiration Date and other license information.

Activate/Retrieve a Device Management License on the M-Series Appliance

Before activating and retrieving a Panorama device management license on the M-Series appliance:

- Register Panorama.
- Locate the auth-codes for the product/subscription you purchased. When you placed your order, Palo Alto Networks Customer Service sent you an email that listed the auth-code associated with the purchase. If you cannot locate this email, contact Customer Support to obtain your codes before proceeding.

After you activate and retrieve the license, the **Panorama > Licenses** page displays the associated issuance date, expiration date, and the number of devices that the license enables Panorama to manage.



To activate and retrieve the license, the options are:

Activate/Retrieve a Device Management License on the M-Series Appliance	
• Use the web interface to activate and retrieve the license. Select this option if Panorama is ready to connect to the Palo Alto Networks update server (you completed the task Perform Initial Configuration of the M-Series Appliance) but you have not activated the license on the Palo Alto Networks Support website.	1. Select **Panorama > Licenses** and click **Activate feature using authorization code**. 2. Enter the **Authorization Code** and click **OK**. Panorama retrieves and activates the license.

Activate/Retrieve a Device Management License on the M-Series Appliance (Continued)	
• Retrieve the license key from the license server. If Panorama is not ready to connect to the update server (for example, you have not completed the initial M-Series appliance setup), you can activate the license on the Support website so that, when Panorama is ready to connect, you can then use the web interface to retrieve the activated license. The process of retrieving an activated license is faster than the process of both retrieving and activating.	1. Activate the license on the Palo Alto Networks Support website. a. On a host with Internet access, access the Palo Alto Support website (https://support.paloaltonetworks.com) in a browser and log in. b. In the **Assets** tab, find your M-Series appliance and, in the Action column, click the edit icon. c. Enter the **Authorization Code** and click **Add** to activate the license. 2. Configure Panorama to connect to the update server: see Perform Initial Configuration of the M-Series Appliance. 3. Select **Panorama > Licenses** and click **Retrieve license keys from the license server**. Panorama retrieves the activated license.
• Manually upload the license from a host to Panorama. Panorama must have access to that host. If Panorama is set up (you completed the task Perform Initial Configuration of the M-Series Appliance) but does not have a connection to the update server, activate the license on the Support website, download it to a host that has a connection to the update server, then upload it to Panorama.	1. Activate and download the license from the Palo Alto Networks Support website. a. On a host with Internet access, access the Palo Alto Support website (https://support.paloaltonetworks.com) in a browser and log in. b. In the **Assets** tab, find your M-Series appliance and, in the Action column, click the edit icon. c. Enter the **Authorization Code** and click **Add** to activate the license. d. In the Action column, click the download icon and save the license key file to the host. 2. In the Panorama web interface, select **Panorama > Licenses**, click **Manually upload license key** and click **Browse**. 3. Select the key file you downloaded to the host and click **Open**. 4. Click **OK** to upload the activated license key.

 © Palo Alto Networks, Inc.

Install Content and Software Updates for Panorama

A valid support subscription enables access to the Panorama software image and release notes. To take advantage of the latest fixes and security enhancements, it is a good idea to upgrade to the latest software update or to the update version that your reseller or a Palo Alto Networks Systems Engineer recommends. The procedure to install software and content updates depends on whether Panorama has a direct connection to the Internet and whether it has a high availability (HA) configuration.

- Panorama, Log Collector, and Firewall Version Compatibility
- Install Updates for Panorama with HA Configuration
- Install Updates for Panorama with Internet Connection
- Install Updates for Panorama without Internet Connection

Panorama, Log Collector, and Firewall Version Compatibility

Palo Alto Networks highly recommends running the same Panorama software version on both the Panorama management server and the Dedicated Log Collectors.

Panorama can manage firewalls running PAN-OS software versions that match the Panorama version or are earlier than the Panorama version. The exception is that Panorama 6.1 and later versions cannot push configurations to firewalls running PAN-OS 6.0.0 through 6.0.3.

The content release version on the Panorama management server must be the same as or higher than the content release versions on the Dedicated Log Collectors and managed firewalls.



Palo Alto Networks recommends installing the same Applications database version on Panorama as on the Dedicated Log Collectors and firewalls.

Regardless of whether your subscriptions include the Applications database or Applications and Threats database, Panorama installs only the Applications database. Panorama and Dedicated Log Collectors do not enforce policy rules and therefore do not need the threat signatures from the Threats database. The Applications database contains threat metadata (such as threat IDs and names) that you use on Panorama and the Dedicated Log Collectors when defining policy rules to push to managed firewalls and when interpreting threat information in logs and reports. Firewalls require the full Applications and Threats database to match the identifiers recorded in logs with the corresponding threat, URL, or application names. Refer to the Release Notes for the minimum content release version you must install for a Panorama release.

Install Updates for Panorama with HA Configuration

To ensure a seamless failover, the active and passive Panorama peers in an HA pair must have the same Panorama version and the same Applications database version. The following example describes how to upgrade an HA pair with an active peer named Primary_A and the passive peer named Secondary_B.

Install Updates for Panorama with HA Configuration		
Step 1	Upgrade the Panorama software version on Secondary_B, the passive peer.	Perform one of the following tasks: • Install Updates for Panorama with Internet Connection • Install Updates for Panorama without Internet Connection After the upgrade, this Panorama transitions to a non-functional state because the software version does not match that of its peer.
Step 2	Suspend Primary_A to trigger a failover.	On Primary_A: 1. Select **Panorama > High Availability**. 2. In the **Operational Commands** section, click **Suspend local Panorama**. 3. Verify that the bottom-right corner of the web interface displays the state as suspended. Upon failover, Secondary_B transitions to an active state.
Step 3	Upgrade the Panorama software version on Primary_A.	Perform one of the following tasks: • Install Updates for Panorama with Internet Connection • Install Updates for Panorama without Internet Connection After rebooting, Primary_A first transitions to the passive state. Then, because preemption is enabled by default, Primary_A automatically transitions to the active state and Secondary_B reverts to the passive state. If you disabled preemption, manually Restore the Primary Panorama to the Active State.
Step 4	Verify that both peers have the Panorama software and content versions.	On the **Dashboard** of each Panorama peer, verify that the Panorama Software Version and Application Version match and that the running configuration is synchronized with the peer.

Install Updates for Panorama with Internet Connection

If Panorama has a direct connection to the Internet, perform the following steps to install content and software updates. If Panorama is deployed in a high availability (HA) configuration, you must upgrade each peer in the order described in Install Updates for Panorama with HA Configuration.



Before upgrading software on an M-Series appliance, we recommend ensuring Panorama has a local Log Collector (**Panorama > Managed Collectors**) and local Collector Group (**Panorama > Collector Groups**).

 © Palo Alto Networks, Inc.

Install Updates for Panorama with Internet Connection		
Step 1	Verify that the updates you plan to install are appropriate for your Panorama deployment.	❑ See Panorama, Log Collector, and Firewall Version Compatibility for critical details about update version compatibility. Refer to the Release Notes for the minimum content version you must install for a Panorama software release. If you will upgrade Log Collectors and firewalls to a particular release, you must first upgrade Panorama to that release. ❑ For a Panorama virtual appliance that runs on an ESXi server, ensure the server meets the requirements listed under Setup Prerequisites for the Panorama Virtual Appliance.
Step 2	Save a backup of the current Panorama configuration file. You can use this backup to restore the configuration if you have problems with the upgrade.  Although Panorama automatically creates a backup of the configuration, the best practice is to create and externally store a backup before upgrading.	1. Log in to Panorama and select **Panorama > Setup > Operations**. 2. Click **Save named Panorama configuration snapshot**, enter a **Name** for the configuration, and click **OK**. 3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, click **OK**, and save the exported file to a location that is external to Panorama.
Step 3	Install the latest content updates.  You must install content updates before software updates. After installing software updates, you can optionally schedule automatic, recurring content updates (Step 6).	1. Select **Panorama > Dynamic Updates** and click **Check Now** to check for the latest updates. If the value in the Action column is **Download**, an update is available. 2. Install the updates in the following sequence. When an installation completes, the Currently Installed column displays a check mark. a. **Download** and **Install** the Applications or Applications and Threats update. Regardless of your subscription, Panorama installs and needs only the Applications content, not the Threats content. For details, see Panorama, Log Collector, and Firewall Version Compatibility. b. **Download** and **Install** any other updates (Antivirus, WildFire, or URL Filtering) one at a time in any sequence.

Install Updates for Panorama with Internet Connection (Continued)		
Step 4	Determine the software upgrade path.	Select **Panorama > Software** and check which version has a check mark in the Currently Installed column. You cannot skip any major Panorama release versions on the path between the current version and the target version. For example, to upgrade from Panorama 5.0.13 to Panorama 7.0.3: 1. Download and install a Panorama 5.1.x release based on your platform: • Panorama virtual appliance—Download and install Panorama 5.1.0 and reboot. • Panorama M-Series appliance: – Download Panorama 5.1.0 without installing or rebooting. – Download and install a later Panorama 5.1.x release and reboot. 2. Download and install Panorama 6.0.0 and reboot. 3. Download and install Panorama 6.1.0 and reboot. 4. Download Panorama 7.0.1 without installing or rebooting (7.0.1 is the base image for the 7.0 release; not 7.0.0). 5. Download and install Panorama 7.0.3 and reboot.
Step 5	Install the software update. Repeat this step for each update you must install based on the upgrade path.	1. Select **Panorama > Software** and click **Check Now** to check for the latest updates. If an update is available, the Action column displays a **Download** link. 2. **Download** the update. When the process finishes, the Action value changes to **Install**. 3. **Install** the update. 4. Reboot Panorama if the update you are installing requires that (see Step 4): • If prompted to reboot, click **Yes**. If you see a `CMS Login` prompt, press Enter without typing a username or password. When the Panorama login prompt appears, enter the username and password you set during initial configuration. • If you are not prompted to reboot, select **Panorama > Setup > Operations** and click **Reboot Panorama** in the Device Operations section.
Step 6	(Best Practice) Schedule recurring, automatic content updates. Panorama does not synchronize content update schedules across HA peers. You must perform this task on both the active and passive Panorama.	1. Select **Panorama > Dynamic Updates**. In the header row for each update type, the **Schedule** is initially set to **None**. Perform the remaining steps for each update type. 2. Click **None** and select the update frequency (**Recurrence**). The frequency options depend on the update type. 3. Select the schedule action: • **Download And Install** (best practice)—Panorama automatically installs updates after downloading them. • **Download Only**—You must manually install updates after Panorama downloads them. 4. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again.

 © Palo Alto Networks, Inc.

Install Updates for Panorama with Internet Connection (Continued)	
Step 7 Configure the Panorama virtual appliance settings on the VMware ESXi server. Required if upgrading from a release earlier than Panorama 5.1 to a Panorama 5.1 or later release running on an ESXi server.	After Panorama reboots, complete the following tasks: 1. Access the VMware vSphere Client and select the **Virtual Machines** tab. 2. Right-click the Panorama virtual appliance and select **Power > Power Off**. 3. Right-click the Panorama virtual appliance and **Edit Settings** as follows: a. Select the **Hardware** tab and allocate **Memory** based on how many firewalls Panorama manages: – 1-10 managed firewalls: 4GB – 11-50 managed firewalls: 8GB – Over 50 managed firewalls: 16GB b. Set the **SCSI Controller** to **LSI Logic Parallel**. c. Select the **Options** tab, select **General Options**, set the **Guest Operating System** to **Linux**, and set the **Version** to **Other Linux (64-bit)**. d. Click **OK**. 4. Right-click the Panorama virtual appliance and select **Power > Power On**.

Install Updates for Panorama without Internet Connection

If Panorama does not have a direct connection to the Internet, perform the following steps to install content and software updates. If Panorama is deployed in a high availability (HA) configuration, you must upgrade each peer in the order described in Install Updates for Panorama with HA Configuration.



Before upgrading software on an M-Series appliance, we recommend ensuring Panorama has a local Log Collector (**Panorama > Managed Collectors**) and local Collector Group (**Panorama > Collector Groups**).

Install Updates for Panorama without Internet Connection	
Step 1 Verify that the updates you plan to install are appropriate for your Panorama deployment.	❑ See Panorama, Log Collector, and Firewall Version Compatibility for critical details about update version compatibility. Refer to the Release Notes for the minimum content release version you must install for a Panorama software release. If you will upgrade Log Collectors and firewalls to a particular release, you must first upgrade Panorama to that release. ❑ For a Panorama virtual appliance that runs on an ESXi server, ensure the server meets the requirements listed under Setup Prerequisites for the Panorama Virtual Appliance.

Install Updates for Panorama without Internet Connection (Continued)		
Step 2	Save a backup of the current Panorama configuration file. You can use this backup to restore the configuration if you have problems with the upgrade. Although Panorama automatically creates a backup of the configuration, the best practice is to create and externally store a backup before upgrading.	1. Log in to Panorama and select **Panorama > Setup > Operations**. 2. Click **Save named Panorama configuration snapshot**, enter a **Name** for the configuration, and click **OK**. 3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, click **OK**, and save the exported file to a location that is external to Panorama.
Step 3	Determine the software upgrade path.	Select **Panorama > Software** and check which version has a check mark in the Currently Installed column. You cannot skip any major Panorama release versions on the path between the current version and the target version. For example, to upgrade from Panorama 5.0.13 to Panorama 7.0.3: 1. Upload and install a Panorama 5.1.x release based on your platform: • Panorama virtual appliance—Upload and install Panorama 5.1.0 and reboot. • Panorama M-Series appliance: – Upload Panorama 5.1.0 without installing or rebooting. – Upload and install a later Panorama 5.1.x release and reboot. 2. Upload and install Panorama 6.0.0 and reboot. 3. Upload and install Panorama 6.1.0 and reboot. 4. Upload Panorama 7.0.1 without installing or rebooting (7.0.1 is the base image for the 7.0 release; not 7.0.0). 5. Upload and install Panorama 7.0.3 and reboot.
Step 4	Download the content and software updates to a host that has Internet access. Panorama must have access to the host.	1. Use a host with Internet access to log in to the Palo Alto Networks Customer Support web site. 2. Download content updates: a. Click **Dynamic Updates** in the Resources section. b. **Download** the desired content update and save the file to the host. Perform this step for each content type you will update. 3. Download software updates: a. Return to the main page of the Palo Alto Networks Customer Support web site and click **Software Updates** in the Resources section. b. Review the Download column to determine the version to install. The filename of the update package indicates the platform: Panorama-ESX-<release> for the Panorama virtual appliance (VMware ESXi server or vCloud Air) or Panorama-m-<release> for the Panorama M-Series appliance. c. Click the filename and save the file to the host.

 © Palo Alto Networks, Inc.

Install Updates for Panorama without Internet Connection (Continued)	
Step 5 Install the content updates. You must install content updates before software updates.	Install the Applications or Applications and Threats update first, and then install any other updates (Antivirus, WildFire, or URL Filtering) one at a time in any sequence. Regardless of whether your subscription includes both Applications and Threats content, Panorama installs and needs only the Applications content. For details, see Panorama, Log Collector, and Firewall Version Compatibility. 1. In Panorama, select **Panorama > Dynamic Updates**. 2. Perform the following steps for each content type: a. Click **Upload**, select the content **Type**, **Browse** to the update, and click **OK**. b. Click **Install From File**, select the **Package Type**, and click **OK**.
Step 6 Install the software update.	1. In the **Panorama > Software** page, click **Upload**. 2. **Browse** to the update, select the **Sync To Peer** check box if Panorama has an HA configuration (to push the software image to the secondary peer), and click **OK**. 3. Click **Install** in the Action column. 4. Reboot Panorama if the update you are installing requires that (see Step 3): • If prompted to reboot, click **Yes**. If you see a `CMS Login` prompt, press Enter without typing a username or password. When the Panorama login prompt appears, enter the username and password you set during initial configuration. • If you are not prompted to reboot, select **Panorama > Setup > Operations** and click **Reboot Panorama** in the Device Operations section.

Install Updates for Panorama without Internet Connection (Continued)	
Step 7 Configure the Panorama virtual appliance settings on the VMware ESXi server. Required if upgrading from a release earlier than Panorama 5.1 to a Panorama 5.1 or later release running on an ESXi server.	After Panorama reboots, complete the following tasks: 1. Access the VMware vSphere Client and select the **Virtual Machines** tab. 2. Right-click the Panorama virtual appliance and select **Power > Power Off**. 3. Right-click the Panorama virtual appliance and **Edit Settings** as follows: a. Select the **Hardware** tab and allocate **Memory** based on how many firewalls Panorama manages: – 1-10 managed firewalls: 4GB – 11-50 managed firewalls: 8GB – Over 50 managed firewalls: 16GB b. Set the **SCSI Controller** to **LSI Logic Parallel**. c. Select the **Options** tab, select **General Options**, set the **Guest Operating System** to **Linux**, and set the **Version** to **Other Linux (64-bit)**. d. Click **OK**. 4. Right-click the Panorama virtual appliance and select **Power > Power On**.

 © Palo Alto Networks, Inc.

Transition to a Different Panorama Platform

When your network requirements change (for example, the logging rate increases), you can migrate the Panorama management server and Dedicated Log Collectors to Panorama Platforms that better support those requirements.

- Migrate from a Panorama Virtual Appliance to an M-Series Appliance
- Migrate from an M-100 Appliance to an M-500 Appliance

Migrate from a Panorama Virtual Appliance to an M-Series Appliance

You can migrate the Panorama configuration from a Panorama virtual appliance to an M-Series appliance. However, you cannot migrate the logs because the log format on the Panorama virtual appliance is incompatible with that on M-Series appliances. Therefore, if you want to maintain access to the old logs stored on the Panorama virtual appliance, you must continue running the Panorama virtual appliance after the migration. The M-Series appliance will collect the new logs that firewalls forward after the migration. After the pre-migration logs expire or become irrelevant due to aging, you can shut down the Panorama virtual appliance.

Migrate from a Panorama Virtual Appliance to an M-Series Appliance	
Step 1 Plan the migration.	❑ Upgrade the software on the Panorama virtual appliance before the migration if the M-Series appliance requires a later release of the current software (the M-500 appliance requires Panorama 7.0 or a later release). For important details about software versions, see Panorama, Log Collector, and Firewall Version Compatibility. ❑ Schedule a maintenance window for the migration. Although firewalls can buffer logs after the Panorama virtual appliance goes offline and then forward the logs after the M-Series appliance comes online, completing the migration during a maintenance window minimizes the risk that logs will exceed the buffer capacities and be lost during the transition between Panorama platforms. ❑ Consider whether to maintain access to the Panorama virtual appliance after the migration to access existing logs. The most efficient approach is to assign a new IP address to the Panorama virtual appliance and reuse its old IP address for the M-Series appliance. This ensures that the Panorama virtual appliance remains accessible and that firewalls can point to the M-Series appliance without you reconfiguring the Panorama IP address on each firewall.

© Palo Alto Networks, Inc.

Migrate from a Panorama Virtual Appliance to an M-Series Appliance (Continued)		
Step 2	Migrate your subscriptions to the new platform.	1. Purchase the new support license and migration license. 2. Provide your sales representative the serial number of the Panorama virtual appliance you will phase out, the auth-code you received when you purchased the new M-Series appliance, and the effective date for the migration. On the effective date, Palo Alto Networks will automatically apply the auth-code to the serial number of the M-Series appliance, phase out support for the Panorama virtual appliance, and trigger support for the M-Series appliance. Consult your sales representative regarding how much time is available to complete the migration after the effective date. At the end of that period, Palo Alto Networks terminates the support entitlement on the Panorama virtual appliance, after which it can no longer receive software or content updates.
Step 3	Export the Panorama configuration from the Panorama virtual appliance.	1. Log in to the Panorama virtual appliance and select **Panorama > Setup > Operations**. 2. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**. 3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. Panorama exports the configuration to your client system as an XML file.
Step 4	Power off the Panorama virtual appliance if you won't need to access to it after the migration or assign a new IP address to its management (MGT) interface if you will need access to it.	To power off the Panorama virtual appliance, see the documentation for your VMware platform. To change the IP address on the Panorama virtual appliance: 1. Select **Panorama > Setup > Management**, and edit the Management Interface Settings. 2. Enter the new **IP Address** and click **OK**. 3. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 5	Perform the initial setup of the M-Series appliance.	1. Rack mount the M-Series appliance. Refer to the M-100 or M-500 Appliance Hardware Reference Guide for instructions. 2. Perform Initial Configuration of the M-Series Appliance to define the network connections required to activate licenses and install updates. 3. Register Panorama. 4. Activate a Panorama Support License. 5. Activate/Retrieve a Device Management License on the M-Series Appliance. Use the auth-code associated with the migration license. 6. Install Content and Software Updates for Panorama. Install the same versions as those on the Panorama virtual appliance.

 © Palo Alto Networks, Inc.

Migrate from a Panorama Virtual Appliance to an M-Series Appliance (Continued)		
Step 6	Load the Panorama configuration snapshot that you exported from the Panorama virtual appliance into the M-Series appliance.	1. On the M-Series appliance, select **Panorama > Setup > Operations**. 2. Click **Import named Panorama configuration snapshot**, **Browse** to the Panorama configuration file you exported from the Panorama virtual appliance, and click **OK**. 3. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, select a **Decryption Key** (the master key for Panorama), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file. 4. If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid.
Step 7	Modify the configuration on the M-Series appliance. Required if the M-Series appliance will use different values than the Panorama virtual appliance. If you will maintain access to the Panorama virtual appliance to access its logs, use a different hostname and IP address for the M-Series appliance.	1. Select **Panorama > Setup > Management**. 2. Edit the General Settings, modify the **Hostname**, and click **OK**. 3. Edit the Management Interface Settings, modify the values as necessary, and click **OK**.
Step 8	Add the default managed collector and Collector Group back to the M-Series appliance. Loading the configuration from the Panorama virtual appliance (Step 6) removes the default managed collector and Collector Group that are predefined on each M-Series appliance.	1. Configure a Managed Collector that is local to the M-Series appliance. 2. Configure a Collector Group for the default managed collector. 3. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 9	Synchronize the M-Series appliance with the firewalls to resume firewall management. Complete this step during a maintenance window to minimize network disruption.	1. On the M-Series appliance, select **Panorama > Managed Devices** and verify that the Device State column displays **Connected** for the firewalls. At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls. 2. Click **Commit**, for the **Commit Type** select **Device Group**, select every device group, select the **Include Device and Network Templates** check box, and click **Commit** again. 3. In the **Panorama > Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

Migrate from an M-100 Appliance to an M-500 Appliance

You can migrate the Panorama configuration and firewall logs from an M-100 appliance to an M-500 appliance in Panorama mode (Panorama management server). You can also migrate the firewall logs from an M-100 appliance to an M-500 appliance in Log Collector mode (Dedicated Log Collector). Because all the Log Collectors in a Collector Group must be the same hardware model, you must migrate all or none of the M-100 appliances in any Collector Group.

In the following procedure, the Panorama management server is deployed in an active/passive high availability (HA) configuration, you will migrate both the configuration and logs, and the M-500 appliances will reuse the IP addresses from the M-100 appliances.



If you will migrate only the logs and not the Panorama configuration, perform the steps under Migrate Logs to a New M-Series Appliance in Log Collector Mode or Migrate Logs to a New M-Series Appliance in Panorama Mode.

Migrate from an M-100 Appliance to an M-500 Appliance	
Step 1 Plan the migration.	❑ Upgrade the software on the M-100 appliance if its current release is earlier than 7.0; the M-500 appliance requires Panorama 7.0 or a later release. For important details about software versions, see Panorama, Log Collector, and Firewall Version Compatibility. ❑ Forward the System and Config logs that Panorama and Log Collectors generate to an external destination before the migration if you want to preserve those logs. The M-Series appliance in Panorama mode stores these log types on its SSD, which you cannot move between platforms. You can move only the RAID drives, which store firewall logs. ❑ Schedule a maintenance window for the migration. Although firewalls can buffer logs after the M-100 appliance goes offline and then forward the logs after the M-500 appliance comes online, completing the migration during a maintenance window minimizes the risk that logs will exceed the buffer capacities and be lost during the transition between Panorama platforms.
Step 2 Migrate your subscriptions to the new platform.	1. Purchase the new support license and migration license. 2. Provide your sales representative the serial number of the M-100 appliance you will phase out, the auth-code you received when you purchased the M-500 appliance, and the effective date for the migration. On the effective date, Palo Alto Networks will automatically apply the auth-code to the serial number of the M-500 appliance, phase out support for the M-100 appliance, and trigger support for the M-500 appliance. Consult your sales representative regarding how much time is available to complete the migration after the effective date. At the end of that period, Palo Alto Networks terminates the support entitlement on the M-100 appliance; it can no longer receive software or content updates.

 © Palo Alto Networks, Inc.

Migrate from an M-100 Appliance to an M-500 Appliance (Continued)		
Step 3	Export the Panorama configuration from each M-100 appliance in Panorama mode.	Perform this task on each M-100 appliance HA peer: 1. Log in to the M-100 appliance and select **Panorama > Setup > Operations**. 2. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**. 3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. Panorama exports the configuration to your client system as an XML file.
Step 4	Power off each M-100 appliance in Panorama mode.	1. Log in to the M-100 appliance HA peer that you will power off. 2. Select **Panorama > Setup > Operations**, and click **Shutdown Panorama**.
Step 5	Perform the initial setup of each M-500 appliance.	1. Rack mount the M-500 appliances. Refer to the M-500 Appliance Hardware Reference Guide for instructions. 2. Perform Initial Configuration of the M-Series Appliance to define the network connections required to activate licenses and install updates. 3. Register Panorama. 4. Activate a Panorama Support License. 5. Activate a device management license. Use the auth-code associated with the migration license. 6. Install Content and Software Updates for Panorama. Install the same versions as those on the M-100 appliance. 7. (Dedicated Log Collector only) Set up the M-Series Appliance as a Log Collector.
Step 6	Load the Panorama configuration snapshot that you exported from each M-100 appliance into each M-500 appliance in Panorama mode (both HA peers).	Perform this task on each M-500 appliance HA peer: 1. Log in to the M-500 appliance and select **Panorama > Setup > Operations**. 2. Click **Import named Panorama configuration snapshot**, **Browse** to the configuration file you exported from the M-100 appliance that has the same HA priority (primary or secondary) as the M-500 appliance will have, and click **OK**. 3. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, select a **Decryption Key** (the master key for Panorama), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file. If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid. 4. Click **Commit** and **Validate Changes**. Resolve any errors before proceeding. 5. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

Migrate from an M-100 Appliance to an M-500 Appliance (Continued)		
Step 7	Synchronize the configuration between the M-500 appliance HA peers in Panorama mode.	1. On the active M-500 appliance, select the **Dashboard** tab and, in the High Availability widget, click **Sync to peer**. 2. In the High Availability widget, verify that the **Local** (primary M-500 appliance) is **active**, the **Peer** is passive, and the **Running Config** is **synchronized**.
Step 8	Move the RAID drives from each M-100 appliance to its replacement M-500 appliance to migrate the logs collected from firewalls.	In the following tasks, skip any steps that you already completed on the M-500 appliance. • Migrate Logs to a New M-Series Appliance in Panorama Mode. Migrate logs from the M-100 appliance only if it uses a default managed collector for log collection. • Migrate Logs to a New M-Series Appliance in Log Collector Mode.
Step 9	Synchronize the active M-500 appliance in Panorama mode with the firewalls to resume firewall management. Complete this step during a maintenance window to minimize network disruption.	1. In the active M-500 appliance, select **Panorama > Managed Devices**, and verify that the Device State column displays **Connected** for the firewalls. At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls. 2. Click **Commit**, for the **Commit Type** select **Device Group**, select every device group, select the **Include Device and Network Templates** check box, and click **Commit** again. 3. In the **Panorama > Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

 © Palo Alto Networks, Inc.

Access and Navigate Panorama Management Interfaces

Panorama provides three management interfaces:

- **Web interface**—The Panorama web interface is purposefully designed with a similar look and feel to the firewall web interface. If you are already familiar with the latter, you can navigate, complete administrative tasks, and generate reports from the Panorama web interface with relative ease. This graphical interface allows you to access Panorama using HTTPS and it is the best way to perform administrative tasks. See Log in to the Panorama Web Interface and Navigate the Panorama Web Interface. If you need to enable HTTP access to Panorama, edit the Management Interface Settings on the **Panorama > Setup > Management** tab.
- **Command line interface (CLI)**—The CLI is a no-frills interface that allows you to type through the commands in rapid succession to complete a series of tasks. The CLI supports two command modes—operational and configuration—and each mode has its own hierarchy of commands and statements. When you get familiar with the nesting structure and the syntax for the commands, the CLI allows quick response times and offers administrative efficiency. See Log in to the Panorama CLI.
- **XML API**—The XML-based API is provided as a web service that is implemented using HTTP/HTTPS requests and responses. It allows you to streamline your operations and integrate with existing, internally developed applications and repositories. For information on how to use the Panorama API interface, refer to the PAN-OS and Panorama XML API Usage Guide.

Log in to the Panorama Web Interface

Log in to the Panorama Web Interface		
Step 1	Log in to the Panorama web interface.	Using a secure connection (https) from a web browser, log in using the IP address and password you assigned during initial configuration (https://<IP address>).
Step 2	(Optional) Enable HTTP and Secure Shell (SSH) access.	1. Select **Panorama > Setup > Management** and edit the Management Interface Settings. 2. Select which management services to allow on the interface. For example, select **HTTP** and **SSH**. 3. Click **OK**.

Navigate the Panorama Web Interface

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

The Panorama web interface includes the following tabs:

Tab	Description
Dashboard	View general information about the Panorama model and network access settings. This tab includes widgets that display information about applications, logs, system resources, and system settings.
ACC	View the overall risk and threat level on the network, based on information that Panorama gathered from the managed firewalls.
Monitor	View and manage logs and reports.
Panorama	Configure Panorama, manage licenses, set up high availability, access software updates and security alerts, manage administrative access, and manage the deployed firewalls and Log Collectors.
Device Groups > Policies	Create centralized policies and apply the configuration to multiple firewalls/device groups. You must Add a Device Group for this tab to display.
Device Groups > Objects	Define policy objects that policy rules can reference and that managed firewalls/device groups can share. You must Add a Device Group for this tab to display.
Templates > Network	Configure network setting, such as network profiles, that can be applied to the managed firewalls. You must Add a Template for this tab to display.
Templates > Device	Configure device configuration, such as server profiles and admin roles, that can be applied to the managed firewalls. You must Add a Template for this tab to display.

Log in to the Panorama CLI

You can log in to the Panorama CLI using a serial port connection or access remotely using a Secure Shell (SSH) client.

Log in to the Panorama CLI	
• Use SSH to log in to the Panorama CLI. The same instructions apply for an M-Series appliance in Log Collector mode.	1. Make sure that you have the following: • A computer with network access to Panorama • Panorama IP address • SSH is enabled on the Management interface. See (Optional) Enable HTTP and Secure Shell (SSH) access. 2. To access the CLI using SSH: a. Enter the Panorama IP address in the SSH client and use port 22. b. Enter your administrative access credentials when prompted. After successfully logging in, the CLI prompt displays in Operational mode. For example: `admin@ABC_Sydney>` See Configure an Administrator with SSH Key-Based Authentication for the CLI.

 © Palo Alto Networks, Inc.

Log in to the Panorama CLI	
• Change to Configuration mode.	To go into Configuration mode, enter the following command at the prompt: `admin@ABC_Sydney>` **`configure`** The prompt changes to `admin@ABC_Sydney#`
• Use a serial port connection to log in to the Panorama CLI.	1. Make sure that you have the following: • A null-modem serial cable that connects Panorama to a computer with a DB-9 serial port • A terminal emulation program running on the computer 2. Use the following settings in the terminal emulation software to connect: 9600 baud; 8 data bits; 1 stop bit; No parity; No hardware flow control. 3. Enter your administrative access credentials when prompted.

Set Up Administrative Access to Panorama

Panorama implements Role-Based Access Control (RBAC) to enable you to specify the privileges and responsibilities of administrators. The following topics describe how to create administrator roles, access domains, and accounts for accessing the Panorama web interface and command line interface (CLI):

- Configure an Admin Role Profile
- Configure an Access Domain
- Configure Administrative Accounts and Authentication

Configure an Admin Role Profile

Admin Role profiles are custom Administrative Roles that enable you to define granular administrative access privileges to ensure protection for sensitive company information and privacy for end users. As a best practice, create Admin Role profiles that allow administrators to access only the areas of the management interfaces required to perform their jobs.

Configure an Admin Role Profile	
Step 1	Select **Panorama > Admin Roles** and click **Add**.
Step 2	Enter a **Name** for the profile and select the **Role** type: **Panorama** or **Device Group and Template**.
Step 3	Configure access privileges to each functional area of Panorama (**Web UI**) and firewalls (**Context Switch UI**) by toggling the icons to the desired setting: Enable (read-write), Read Only, or Disable.  If administrators with custom roles will commit device group or template changes to managed firewalls, you must give those roles read-write access to **Panorama > Device Groups** and **Panorama > Templates**. If you upgrade from an earlier Panorama version, the upgrade process provides read-only access to those nodes. You cannot manage access to the firewall CLI or XML API through context-switching privileges in Panorama roles.
Step 4	If the **Role** type is **Panorama**, configure access to the **XML API** by toggling the Enabled/Disabled icon for each functional area.
Step 5	If the **Role** type is **Panorama**, select an access level for the **Command Line** interface: **None** (default), **superuser**, **superreader**, or **panorama-admin**.
Step 6	Click **OK** to save the profile.

Configure an Access Domain

Use Access Domains to define access for Device Group and Template administrators for specific device groups and templates, and also to control the ability of those administrators to switch context to the web interface of managed firewalls. Panorama supports up to 4,000 access domains.

 © Palo Alto Networks, Inc.

Configure an Access Domain	
Step 1	Select **Panorama > Access Domain** and click **Add**.
Step 2	Enter a **Name** to identify the access domain.
Step 3	Select an access privilege for **Shared Objects**: • **write**—Administrators can perform all operations on Shared objects. This is the default value. • **read**—Administrators can display and clone but cannot perform other operations on Shared objects. When adding non-Shared objects or cloning Shared objects, the destination must be a device group within the access domain, not the Shared location. • **shared-only**—Administrators can add objects only to the Shared location. Administrators can display, edit, and delete Shared objects but cannot move or clone them. A consequence of this option is that administrators can't perform any operations on non-Shared objects other than to display them. An example of why you might select this option is for an organization that requires all objects to be in a single, global repository.
Step 4	Toggle the icons in the **Device Groups** tab to enable read-write or read-only access for device groups in the access domain. If you set the **Shared Objects** access to **shared-only**, Panorama applies read-only access to the objects in any device groups for which you specify read-write access.
Step 5	Select the **Templates** tab and **Add** each template you want to assign to the access domain.
Step 6	Select the **Device Context** tab, select firewalls to assign to the access domain, and click **OK**. Administrators can access the web interface of these firewalls by using the **Context** drop-down in Panorama.

Configure Administrative Accounts and Authentication

If you have already configured Administrative Roles, external authentication services (if applicable), and Access Domains (for Device Group and Template administrators), you can Configure an Administrative Account. Otherwise, perform one of the other procedures listed below to configure administrative accounts for specific types of authentication.

- Configure an Administrative Account
- Configure an Administrator with Kerberos SSO, External, or Local Authentication
- Configure an Administrator with Certificate-Based Authentication for the Web Interface
- Configure an Administrator with SSH Key-Based Authentication for the CLI
- Configure RADIUS Vendor-Specific Attributes for Administrator Authentication

Configure an Administrative Account

Administrative accounts specify Administrative Roles, Administrative Authentication methods, and Access Domains for Panorama administrators.

You can't add an administrator account to a Dedicated Log Collector (M-Series appliance in Log Collector mode). Only the predefined administrator account with the default username (admin) is available on Dedicated Log Collectors.

Configure an Administrative Account	
Step 1	Select **Panorama > Administrators** and click **Add**.
Step 2	Enter a user **Name** for the administrator.
Step 3	Select an **Authentication Profile** if Panorama will use Kerberos SSO or an external service for authentication. If Panorama will use local authentication, set the **Authentication Profile** to **None** and enter a **Password** and then **Confirm Password**.
Step 4	For the **Administrator Type**, select the role type: • **Dynamic**—Select a predefined administrator role. • **Custom Panorama Admin**—Select the Admin Role **Profile** you created for this administrator. • **Device Group and Template Admin**—In the Access Domain to Administrator Role section, add access domains and map each one to an Admin Role profile: 1. Click **Add**. 2. Select an Access Domain from the drop-down. 3. Click the adjacent Admin Role cell. 4. Select an Admin Role profile.
Step 5	(Optional) Select a **Password Profile** if you will use local authentication.
Step 6	Click **OK** and **Commit**, select **Panorama** for the **Commit Type**, and click **Commit** again.

Configure an Administrator with Kerberos SSO, External, or Local Authentication

When you configure Administrative Authentication for an administrator account, you can combine Kerberos single sign-on (SSO) authentication with an external authentication service or with local authentication. You can also configure the administrator to use only one of those authentication methods.

Configure an Administrator with Kerberos SSO, External, or Local Authentication		
Step 1	Create a Kerberos keytab. Required for Kerberos SSO authentication.	Create a Kerberos keytab. A keytab is a file that contains Kerberos account information (principal name and hashed password) for Panorama.
Step 2	Configure access domains. Required for Device Group and Template administrators.	Configure an Access Domain.
Step 3	Configure Admin Role profiles. Required if you are assigning a custom role to the administrator.	Configure an Admin Role Profile.

 © Palo Alto Networks, Inc.

Configure an Administrator with Kerberos SSO, External, or Local Authentication (Continued)		
Step 4	Configure access to an external authentication service if you will use one.	Select **Panorama > Server Profiles**, select the authentication service type (**RADIUS**, **TACACS+**, **LDAP**, or **Kerberos**), and configure the server profile: • Configure a RADIUS Server Profile. • Configure a TACACS+ Server Profile. • Configure an LDAP Server Profile. • Configure a Kerberos Server Profile.
Step 5	Configure an authentication profile. Required for Kerberos SSO or external authentication. If your administrators are in multiple Kerberos realms, you can create an authentication profile for each realm and assign all the profiles to an authentication sequence. You can then assign the same authentication sequence to all administrators. For details, see Authentication Profiles and Sequences.	Configure an authentication profile or sequence.
Step 6	Configure an administrator.	Configure an Administrative Account.

Configure an Administrator with Certificate-Based Authentication for the Web Interface

As a more secure alternative to password-based authentication to the Panorama web interface, you can configure certificate-based authentication for administrator accounts that are local to Panorama. Certificate-based authentication involves the exchange and verification of a digital signature instead of a password.

Configuring certificate-based authentication for any administrator disables the username/password logins for all administrators on Panorama and all administrators thereafter require the certificate to log in.

Configure an Administrator with Certificate-Based Authentication for the Web Interface		
Step 1	Generate a certificate authority (CA) certificate on Panorama. You will use this CA certificate to sign the client certificate of each administrator.	Create a self-signed root CA certificate. Alternatively, you can import a certificate from your enterprise CA.

Configure an Administrator with Certificate-Based Authentication for the Web Interface (Continued)		
Step 2	Configure a certificate profile for securing access to the web interface.	1. Select **Panorama > Certificate Management > Certificate Profile** and click **Add**. 2. Enter a **Name** for the certificate profile and set the **Username Field** to **Subject**.. 3. Select **Add** in the CA Certificates section and select the **CA Certificate** you just created. 4. Click **OK** to save the profile.
Step 3	Configure Panorama to use the certificate profile for authenticating administrators.	1. Select the **Panorama > Setup > Management** and edit the Authentication Settings. 2. Select the **Certificate Profile** you just created and click **OK**.
Step 4	Configure the administrator accounts to use client certificate authentication.	Configure an Administrative Account for each administrator who will access the Panorama web interface. Select the **Use only client certificate authentication (Web)** check box. If you have already deployed client certificates that your enterprise CA generated, skip to Step 8. Otherwise, continue with Step 5.
Step 5	Generate a client certificate for each administrator.	Generate a certificate on Panorama. In the **Signed By** drop-down, select the CA certificate you created.
Step 6	Export the client certificates.	1. Export the certificates. 2. **Commit** your changes. Panorama restarts and terminates your login session. Thereafter, administrators can access the web interface only from client systems that have the client certificate you generated.
Step 7	Import the client certificate into the client system of each administrator who will access the web interface.	Refer to your web browser documentation as needed to complete this step.
Step 8	Verify that administrators can access the web interface.	1. Open the Panorama IP address in a browser on the computer that has the client certificate. 2. When prompted, select the certificate you imported and click **OK**. The browser displays a certificate warning. 3. Add the certificate to the browser exception list. 4. Click **Login**. The web interface should appear without prompting you for a username or password.

Configure an Administrator with SSH Key-Based Authentication for the CLI

For administrators who use Secure Shell (SSH) to access the Panorama CLI, SSH keys provide a more secure authentication method than passwords. SSH keys almost eliminate the risk of brute-force attacks, provide the option for two-factor authentication (private key and passphrase), and don't send passwords over the network. SSH keys also enable automated scripts to access the CLI.

 © Palo Alto Networks, Inc.

Configure an Administrator with SSH Key-Based Authentication for the CLI		
Step 1	Use an SSH key generation tool to create an asymmetric key pair on the client system of the administrator. The supported key formats are IETF SECSH and Open SSH. The supported algorithms are DSA (1024 bits) and RSA (768-4096 bits).	For the commands to generate the key pair, refer to your SSH client documentation. The public key and private key are separate files. Save both to a location that Panorama can access. For added security, enter a passphrase to encrypt the private key. Panorama prompts the administrator for this passphrase during login.
Step 2	Configure the administrator account to use public key authentication.	1. Configure an Administrative Account. • Configure one of two authentication methods to use as a fallback if SSH key authentication fails: – Select an **Authentication Profile** if you configured one for the administrator. – Select **None** and enter a **Password** and **Confirm Password**. • Select the **Use Public Key Authentication (SSH)** check box, click **Import Key**, **Browse** to the public key you just generated, and click **OK**. 2. Click **OK** and **Commit**, select **Panorama** for the **Commit Type**, and click **Commit** again.
Step 3	Configure the SSH client to use the private key to authenticate to Panorama.	Perform this task on the client system of the administrator. Refer to your SSH client documentation as needed to complete this step.
Step 4	Verify that the administrator can access the Panorama CLI using SSH key authentication.	1. Use a browser on the client system of the administrator to go to the Panorama IP address. 2. Log in to the Panorama CLI as the administrator. After entering a username, you will see the following output (the key value is an example): `Authenticating with public key "dsa-key-20130415"` 3. If prompted, enter the passphrase you defined when creating the keys.

Configure RADIUS Vendor-Specific Attributes for Administrator Authentication

The following procedure provides an overview of the tasks required to configure RADIUS Vendor-Specific Attributes (VSAs) for administrator authentication to Panorama. For detailed instructions, refer to the following Knowledge Base (KB) articles:

- For Windows 2003 Server, Windows 2008 (and later), and Cisco ACS 4.0—RADIUS Vendor-Specific Attributes (VSAs).
- For Cisco ACS 5.2—Configuring Cisco ACS 5.2 for use with Palo Alto VSA.

Be sure to complete the following two tasks before you start this procedure:

- ❒ Create the administrative accounts in the directory service that your network uses (for example, Active Directory).
- ❒ Set up a RADIUS server that can communicate with that directory service.

Use RADIUS Vendor-Specific Attributes for Account Authentication	
Step 1 Configure Panorama.	1. Configure an Admin Role Profile if the administrator will use a custom role. 2. Configure an Access Domain if the administrator will use a Device Group and Template role. 3. Configure a RADIUS server profile. 4. Configure an authentication profile. Set the authentication **Type** to **RADIUS** and assign the RADIUS **Server Profile**. 5. Configure Panorama to use the authentication profile for administrator access: select **Panorama > Setup > Management**, edit the Authentication Settings, and select the **Authentication Profile**. 6. Click **OK** and **Commit**, select **Panorama** for the **Commit Type**, and click **Commit** again.
Step 2 Configure the RADIUS server.	1. Add the Panorama IP address or hostname as the RADIUS client. 2. Define the VSAs for administrator authentication. You must specify the vendor code (25461 for Panorama) and the VSA name, number, and value.

 © Palo Alto Networks, Inc.

Manage Firewalls

To use Panorama for managing Palo Alto Networks firewalls, you must add the firewalls as managed devices and then assign them to device groups and to templates or template stacks. The following tasks best suit a first-time firewall deployment. Before proceeding, review Plan Your Deployment to understand the deployment options.

- Add a Firewall as a Managed Device
- Manage Device Groups
- Manage Templates and Template Stacks
- Transition a Firewall to Panorama Management
- Use Case: Configure Firewalls Using Panorama



To view the **Objects** and **Policies** tabs on the Panorama web interface, you must first create at least one device group. To view the **Network** and **Device** tabs, you must create at least one template. These tabs contain the options by which you configure and manage the firewalls on your network.

Add a Firewall as a Managed Device

To use Panorama for central management of firewalls, the first step is to add them as managed devices. Before starting, collect the firewall serial numbers and prepare each firewall as follows:

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

You can then add the firewalls as managed devices on Panorama.



When you add a firewall as a managed device, it uses an SSL connection with AES-256 encryption to register with Panorama. Panorama and the firewall authenticate each other using 2,048-bit certificates and use the SSL connection for configuration management and log collection.

Add a Firewall as a Managed Device	
Step 1 Add the firewall to Panorama.	1. Select **Panorama > Managed Devices** and click **Add**. 2. Enter the serial number for each firewall (one entry per line) that you want to manage centrally using Panorama, and then click **OK**. The Managed Devices page displays the new device. 3. (Optional) Add a **Tag**. Tags make it easier for you to find a device from a large list; they help you to dynamically filter and refine the list of firewalls that display. For example, if you add a tag called branch office, you can filter for all branch office devices across your network. a. Select the check box beside the managed device and click **Tag**. b. Click **Add**, enter a string of up to 31 characters (no empty spaces), and click **OK**. 4. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 2 Verify that the firewall is connected to Panorama.	In the **Panorama > Managed Devices** page, the Device State column displays whether the firewall is connected or disconnected to Panorama.

 © Palo Alto Networks, Inc.

Manage Device Groups

- Add a Device Group
- Create a Device Group Hierarchy
- Create Objects for Use in Shared or Device Group Policy
- Revert to Inherited Object Values
- Manage Unused Shared Objects
- Manage Precedence of Inherited Objects
- Move or Clone a Policy Rule or Object to a Different Device Group
- Select a URL Filtering Vendor on Panorama
- Push a Policy Rule to a Subset of Firewalls
- Manage the Rule Hierarchy

Add a Device Group

After adding firewalls (see Add a Firewall as a Managed Device), you can group them into Device Groups (up to 256), as follows. Be sure to assign both firewalls in an active-passive high availability (HA) configuration to the same device group so that Panorama will push the same policy rules and objects to those firewalls. PAN-OS doesn't synchronize pushed rules across HA peers. To manage rules and objects at different administrative levels in your organization, Create a Device Group Hierarchy.

Add a Device Group	
Step 1	Select **Panorama > Device Groups**, and click **Add**.
Step 2	Enter a unique **Name** and a **Description** to identify the device group.
Step 3	In the Devices section, select check boxes to assign firewalls to the group. To search a long list of firewalls, use the Filters. You can assign any firewall to only one device group. You can assign each virtual system on a firewall to a different device group.
Step 4	(Optional) Select the **Group HA Peers** check box for firewalls that are HA peers. The firewall name of the passive or active-secondary peer is in parentheses.
Step 5	Select the **Parent Device Group** (default is **Shared**) that will be just above the device group you are creating in the device group hierarchy.
Step 6	If your policy rules will reference users and groups, assign a **Master** firewall. This will be the only firewall in the device group from which Panorama gathers username and user group information.
Step 7	Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 8	Click **Commit**, for the **Commit Type** select **Device Group**, select the device group you just created, and click **Commit** again.

Create a Device Group Hierarchy

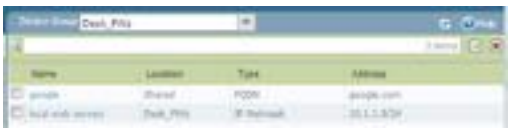
Create a Device Group Hierarchy	
Step 1 Plan the Device Group Hierarchy.	1. Decide the device group levels, and which firewalls and virtual systems you will assign to each device group and the Shared location. You can assign any one firewall or virtual system (vsys) to only one device group. If a device group will be just an organizational container for lower level device groups, you don't need to assign devices to it. 2. Remove firewall or vsys assignments from existing device groups if those assignments don't fit your planned hierarchy. a. Select **Panorama > Device Groups** and select the device group. b. In the Devices section, clear the check boxes of firewalls and virtual systems you want to remove, and click **OK**. 3. If necessary, add more firewalls that you will assign to device groups: see Add a Firewall as a Managed Device.
Step 2 For each top-level device group, Add a Device Group.	1. In the **Panorama > Device Groups** page, click **Add** and enter a **Name** to identify the device group. 2. In the Devices section, select check boxes to assign devices and virtual systems to the device group. 3. Leave the **Parent Device Group** option at **Shared** (the default) and click **OK**.
Step 3 For each lower-level device group, Add a Device Group.	• For new device groups at each lower level, repeat Step 2 but set the **Parent Device Group** to a device group at the next level above. • For each existing device group, in the **Device Groups** page, select the device group to edit it, select a **Parent Device Group**, and click **OK**. If you move a device group to a different parent, all its descendant device groups move with it, along with all devices, policy rules, and objects associated with the device group and its descendants. If the new parent is in another access domain, the moved device group will no longer have membership in the original access domain. If the new access domain has read-write access for the parent device group, it will also have read-write access for the moved device group. If the new access domain has read-only access for the parent, it will have no access for the moved device group. To reconfigure access for device groups, see Configure an Access Domain.

 © Palo Alto Networks, Inc.

Create a Device Group Hierarchy (Continued)		
Step 4	Configure, move, and clone objects and policy rules as needed to account for inheritance in the device group hierarchy.	• Create Objects for Use in Shared or Device Group Policy, or edit existing objects. You can edit objects only at their *location*: the device group to which they are assigned. Descendant device groups inherit read-only instances of the objects from that location. However, you can optionally Override inherited object values. • Create or edit policies. • Move or Clone a Policy Rule or Object to a Different Device Group.
Step 5	Override inherited object values. Applicable only if object values in a particular device group must differ from the values inherited from an ancestor device group. After overriding an object, you can override it again in descendant device groups. However, you can never override shared or predefined (default) objects. In the **Objects** tab, inherited objects have a green icon in the Name column, and the Location column displays the ancestor device group.	1. In the **Objects** tab, select the object type (for example, **Objects > Addresses**). 2. Select the **Device Group** that will have the override instance. 3. Select the object and click **Override**. 4. Edit the values. You can't edit the **Name** or **Shared** settings. 5. Click **OK**. The Name column displays a yellow-overlapping-green icon for the object to indicate it is overridden. If necessary, you can later Revert to Inherited Object Values.
Step 6	Save and commit your changes. Perform a Panorama and device group commit after any change to the hierarchy. You must also perform a template commit if a template references objects in a device group (for example, interfaces referencing addresses), and a firewall assigned to the template is no longer assigned to that device group because of a hierarchy change.	1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 2. Click **Commit**, for the **Commit Type** select **Device Group**, select all the device groups you added or changed, and click **Commit** again.

Create Objects for Use in Shared or Device Group Policy

You can use an object in any policy rule that is in the Shared location, or in the same device group as the object, or in descendants of that device group (for details, see Device Group Objects).

Create Objects for Use in Shared or Device Group Policy	
• Create a shared object. In this example, we add a shared object for URL Filtering categories for which we want to trigger alerts.	1. Select the **Objects > Security Profiles > URL Filtering** tab and click **Add.** The **Objects** tab appears only after you Add a Device Group (at least one). 2. Enter a **Name** and a **Description**. 3. Select the **Shared** check box. 4. The **Disable Override** check box is cleared by default, which means you can override inherited instances of the object in all device groups. To disable overrides for the object, select the check box. 5. In the **Categories** tab, select the check box of every Category for which you want notification. 6. In the **Action** column, select **Alert**. 7. Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
• Create a device group object. In this example, we add an address object for specific web servers on your network.	1. Select **Objects > Addresses** and select the **Device Group** in which you will use the object. 2. Click **Add** and enter a **Name** to identify the object. 3. Be sure to leave the **Shared** check box cleared. 4. The **Disable Override** check box is cleared by default, which means you can override inherited instances of the object in device groups that are descendants of the selected **Device Group**. To disable overrides for the object, select the check box. 5. Select the **Type** of address object and the associated value. For example, select **IP Range** and enter the IP address range for the web servers. 6. Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 7. Click **Commit**, for the **Commit Type** select **Device Group**, select the device group to which you added the object, and click **Commit** again.
• View shared objects and device group objects in Panorama. To demonstrate the difference between a shared object and a device group object, the following screenshot includes a shared address object that was created on Panorama. 	In the pages of the **Objects** tab, the Location column indicates whether an object is shared or is specific to a device group. 1. In the **Objects** tab, select the object type (**Objects > Addresses**, in this example). 2. Select the **Device Group** to which you added the object. The **Objects** tab only displays objects that are in the selected **Device Group** or are inherited from an ancestor device group or the Shared location. 3. Verify that the device group object appears. Note that the device group name in the Location column matches the selection in the **Device Group** drop-down.

 © Palo Alto Networks, Inc.

Revert to Inherited Object Values

After overriding the values that a device group object inherits from an ancestor device group, you can revert the object to its ancestor values at any time. In the **Objects** tab, overridden objects have a yellow-overlapping-green icon in the Name column.

If you want to push ancestor values to all overridden objects instead of reverting a specific object, see Manage Precedence of Inherited Objects.
For the steps to override values, see Override inherited object values.
For details on object inheritance and overrides, see Device Group Objects.

Revert an Overridden Object	
Step 1	In the **Objects** tab, select the object type (for example, **Objects > Addresses**) and select the **Device Group** that has an override instance of the object.
Step 2	Select the object, click **Revert**, and click **Yes**. The Name column displays a green icon for the object, indicating that it now inherits all values from an ancestor device group.
Step 3	Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 4	Click **Commit**, for the **Commit Type** select **Device Group**, select the device group in which you reverted the object, and click **Commit** again.

Manage Unused Shared Objects

When you commit Device Groups, by default Panorama pushes all shared objects to firewalls whether or not any shared or device group policy rules reference the objects. However, you can configure Panorama to push only the shared objects that rules reference in the device groups you commit. The **Share Unused Address and Service Objects with Devices** check box enables you to limit the objects that Panorama pushes to the managed firewalls.

On lower-end platforms, such as the PA-200, consider pushing only the relevant shared objects to the managed devices. This is because the number of objects that can be stored on the lower-end platforms is considerably lower than that of the mid- to high-end platforms. Also, if you have many address and service objects that are unused, clearing the **Share Unused Address and Service Objects with Devices** check box reduces the commit times significantly on the devices because the configuration pushed to each device is smaller. Disabling this option may, however, increase the commit time on Panorama. This is because Panorama has to dynamically check whether policy rules reference a particular object.

Manage Unused Shared Objects	
Step 1	Select **Panorama > Setup > Management**, and edit the Panorama Settings.
Step 2	Clear the **Share Unused Address and Service Objects with Devices** check box to push only the shared objects that rules reference, or select the check box to re-enable pushing all shared objects.
Step 3	Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

Manage Precedence of Inherited Objects

By default, when device groups at different levels in the Device Group Hierarchy have an object with the same name but different values (because of overrides, as an example), policy rules in a descendant device group use the object values in that descendant instead of using object values inherited from ancestor device groups. Optionally, you can reverse this order of precedence to push values from the highest ancestor containing the object to all descendant device groups. After you enable this option, the next device group commit replaces any overridden objects in the descendant device groups with the inherited objects.



If a firewall has locally defined objects with the same name as shared or device group objects that Panorama pushes, a commit failure occurs.



If you want to revert a specific overridden object to its ancestor values instead of pushing ancestor values to all overridden objects, see Revert to Inherited Object Values.

Manage Precedence of Inherited Objects	
Step 1	Select **Panorama > Setup > Management** and edit the Panorama Settings.
Step 2	If you want to reverse the default order of precedence, select the **Objects defined in ancestors will take higher precedence** check box. The dialog then displays the **Find Overridden Objects** link, which provides the option to see how many overridden (shadowed) objects will have ancestor values after you commit this change. You can hover over the quantity message to display the object names. If you want to revert to the default order of precedence, clear the check box.
Step 3	Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 4	(Optional) If you selected the **Objects defined in ancestors will take higher precedence** check box, Panorama won't push the ancestor objects until you perform a device group commit: click **Commit**, for the **Commit Type** select **Device Group**, select the desired device groups, and click **Commit** again.

Move or Clone a Policy Rule or Object to a Different Device Group

On Panorama, if a policy rule or object that you will move or clone from a device group has references to objects that are not available in the target device group (**Destination**), you must move or clone the referenced objects and the referencing rule or object in the same operation. In a Device Group Hierarchy, remember that referenced objects might be available through inheritance. For example, shared objects are available in all device groups. You can perform a global find to check for references. If you move or clone an overridden object, be sure that overrides are enabled for that object in the parent device group of the **Destination** (see Create Objects for Use in Shared or Device Group Policy).

Move or Clone a Policy Rule or Object to a Device Group	
Step 1	Log in to Panorama and select the rulebase (for example, **Policy > Security > Pre Rules**) or object type (for example, **Objects > Addresses**).
Step 2	Select the **Device Group** and select one or more rules or objects.

 © Palo Alto Networks, Inc.

Move or Clone a Policy Rule or Object to a Device Group (Continued)	
Step 3	Perform one of the following steps: • (Rules only) **Move > Move to other device group** • (Objects only) **Move** • (Rules or objects) **Clone**
Step 4	In the **Destination** drop-down, select the new device group or **Shared**. The default is the **Device Group** selected in Step 2.
Step 5	(Rules only) Select the **Rule order**: • **Move top** (default)—The rule will come before all other rules. • **Move bottom**—The rule will come after all other rules. • **Before rule**—In the adjacent drop-down, select the rule that comes after the Selected Rules. • **After rule**—In the adjacent drop-down, select the rule that comes before the Selected Rules.
Step 6	The **Error out on first detected error in validation** check box is selected by default, which means Panorama will display the first error it finds and stop checking for more errors. For example, an error occurs if the **Destination** device group doesn't have an object that is referenced in the rule you are moving. When you move or clone many items at once, selecting this check box can simplify troubleshooting. If you clear the check box, Panorama will find all the errors before displaying them. Regardless of this setting, Panorama won't move or clone anything until you fix all the errors for all the selected items.
Step 7	Click **OK** to start the error validation. If Panorama finds errors, fix them and retry the move or clone operation. If Panorama doesn't find errors, it performs the operation.
Step 8	Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 9	Click **Commit**, for the **Commit Type** select **Device Group**, select the original and destination device groups, and click **Commit** again.

Select a URL Filtering Vendor on Panorama

URL filtering enables firewalls to monitor and control web access for your users. The policy rules that you configure to control web access (Security, QoS, Captive Portal, and Decryption rules) reference URL categories. The URL filtering vendor you select on Panorama determines which URL categories are available for referencing in the rules that you add to device groups and push to firewalls.

By default, Panorama uses PAN-DB, a URL filtering database that is tightly integrated into PAN-OS and the Palo Alto Networks threat intelligence cloud. PAN-DB provides high-performance local caching to maximize in-line performance for URL lookups. The other vendor option is BrightCloud, a third-party URL database.



Unlike firewalls, Panorama does not download the URL database and does not require a URL filtering license.

The following topics describe how to change the URL filtering vendor on Panorama or on both Panorama and managed firewalls. You can also change the URL filtering vendor on just the firewalls.

- Must Panorama and Firewalls Have Matching URL Filtering Vendors?
- Change the URL Filtering Vendor on HA Panorama
- Change the URL Filtering Vendor on non-HA Panorama

- Migrate Panorama and HA Firewalls from BrightCloud to PAN-DB
- Migrate Panorama and non-HA Firewalls from BrightCloud to PAN-DB

Must Panorama and Firewalls Have Matching URL Filtering Vendors?

On any single Panorama management server or firewall, only one URL filtering vendor can be active: PAN-DB or BrightCloud. When selecting a vendor for Panorama, you must consider the vendor and PAN-OS version of the managed firewalls:

- PAN-OS 5.0.x and earlier versions—Panorama and the firewalls require matching URL filtering vendors.
- PAN-OS 6.0 or later versions—Panorama and the firewalls do not require matching URL filtering vendors. If a vendor mismatch is detected, the firewall maps the URL categories in the URL Filtering profiles and rules that it received from Panorama to URL categories that align with those of the vendor enabled on the firewall.

Therefore, for a deployment in which some firewalls run PAN-OS 6.0 or later and some firewalls run earlier PAN-OS versions, Panorama must use the same URL filtering vendor as the firewalls that run earlier PAN-OS versions. For example, if firewalls that run PAN-OS 5.0 use PAN-DB, and firewalls that run PAN-OS 7.0 use BrightCloud, Panorama must use PAN-DB.

Change the URL Filtering Vendor on HA Panorama

In a high availability (HA) deployment, each Panorama peer must be in a non-functional state when you change the URL filtering vendor. Therefore, to avoid disrupting Panorama operations, change the URL filtering vendor on the passive Panorama (Panorama2 in this example) and then trigger failover before changing the vendor on the active Panorama (Panorama1 in this example).

Change URL Filtering Vendor on HA Panorama		
Step 1	Change the URL filtering vendor on each Panorama HA peer.  Complete this task on Panorama2 (passive peer) before Panorama1 (active peer).	1. Log in to the Panorama web interface. 2. Select **Panorama > High Availability** and **Suspend local Panorama**. When you perform this step on Panorama1, failover occurs and Panorama2 becomes active. 3. Select **Panorama > Setup > Management** and edit the General Settings. 4. Select the **URL Filtering Database** vendor: **paloaltonetworks** (PAN-DB) or **brightcloud**. 5. Select **Panorama > High Availability** and **Make local Panorama functional**. When you perform this step on Panorama1 with preemption enabled on both HA peers, Panorama1 automatically reverts to active status and Panorama2 reverts to passive status.
Step 2	Verify that the URL categories are available for referencing in policies.	1. Select **Objects > Security Profiles > URL Filtering.** 2. Click **Add** and verify that the **Categories** tab of the URL Filtering profile dialog displays the URL categories associated with the selected vendor.

 © Palo Alto Networks, Inc.

Change the URL Filtering Vendor on non-HA Panorama

Perform this procedure to change the URL filtering vendor on a Panorama management server that is not deployed in a high availability (HA) configuration.

Change URL Filtering Vendor on non-HA Panorama		
Step 1	Change the URL filtering vendor.	1. Select **Panorama > Setup > Management** and edit the General Settings. 2. Select the **URL Filtering Database** vendor: **paloaltonetworks** (PAN-DB) or **brightcloud**.
Step 2	Verify that the URL categories are available for referencing in policies.	1. Select **Objects > Security Profiles > URL Filtering.** 2. Click **Add** and verify that the **Categories** tab of the URL Filtering profile dialog displays the URL categories associated with the selected vendor.

Migrate Panorama and HA Firewalls from BrightCloud to PAN-DB

Perform this procedure to migrate the URL filtering vendor from BrightCloud to PAN-DB on Panorama and firewalls when the firewalls are deployed in a high availability (HA) configuration. In this example, the active (or active-primary) firewall is named fw1 and the passive (or active-secondary) firewall is named fw2. The migration automatically maps BrightCloud URL categories to PAN-DB URL categories.

Migrate Panorama and HA Firewalls from BrightCloud to PAN-DB		
Step 1	Determine which firewalls require new PAN-DB URL filtering licenses.	1. Log in to Panorama and select **Panorama > Device Deployment > Licenses**. 2. Check the URL column to determine which firewalls have PAN-DB licenses and whether the licenses are valid or expired. A firewall can have valid licenses for both BrightCloud and PAN-DB, but only one license can be active.  If you're not sure whether a PAN-DB URL filtering license is active, access the firewall web interface, select **Device > Licenses**, and verify that the **Active** field displays **Yes** in the PAN-DB URL Filtering section. 3. Purchase a new license for each firewall that does not have a valid PAN-DB license. In HA deployments, each firewall peer needs a distinct PAN-DB license and authorization code. Palo Alto Networks sends an email containing activation codes for the licenses you purchase. If you can't find this email, contact Customer Support before proceeding.
Step 2	Change the URL filtering vendor to PAN-DB on Panorama.	Access the Panorama web interface and perform one of the following tasks: • Change the URL Filtering Vendor on HA Panorama • Change the URL Filtering Vendor on non-HA Panorama

Migrate Panorama and HA Firewalls from BrightCloud to PAN-DB (Continued)		
Step 3	Configure the TCP session settings on both firewall HA peers to ensure sessions that are not yet synchronized will fail over when you suspend a peer.	Log in to the CLI of each firewall and run the following command: `> set session tcp-reject-non-syn no`
Step 4	Migrate the URL filtering vendor to PAN-DB on each firewall HA peer. Complete this task on fw2 (passive or active-secondary peer) before fw1 (active or active-primary peer).	1. Access the firewall web interface, select **Device > High Availability > Operational Commands**, and **Suspend local device**. Performing this step on fw1 triggers failover to fw2. 2. Select **Device > Licenses**. 3. In the License Management section, select **Activate feature using authorization code**, enter the **Authorization Code** and click **OK**. Activating the PAN-DB license automatically deactivates the BrightCloud license. 4. In the PAN-DB URL Filtering section, **Download** the seed file, select your region, and click **OK**. 5. Access the Panorama web interface, click **Commit**, set the **Commit Type** to **Device Group**, select the firewall, and click **Commit** again. 6. Access the firewall web interface, select **Device > High Availability > Operational Commands**, and **Make local device functional**. When you perform this step on fw1 with preemption enabled on both firewalls, fw1 automatically reverts to active (or active-primary) status and fw2 reverts to passive (or active-secondary) status.
Step 5	Revert both firewall HA peers to the original TCP session settings.	Run the following command at the CLI of each firewall: `> set session tcp-reject-non-syn yes`

Migrate Panorama and non-HA Firewalls from BrightCloud to PAN-DB

Perform this procedure to migrate the URL filtering vendor from BrightCloud to PAN-DB on Panorama and firewalls when the firewalls are not deployed in a high availability (HA) configuration. The migration automatically maps BrightCloud URL categories to PAN-DB URL categories.

 © Palo Alto Networks, Inc.

Migrate Panorama and non-HA Firewalls from BrightCloud to PAN-DB		
Step 1	Determine which firewalls require new PAN-DB URL filtering licenses.	1. Log in to Panorama and select **Panorama > Device Deployment > Licenses**. 2. Check the URL column to determine which firewalls have PAN-DB licenses and whether the licenses are valid or expired. A firewall can have valid licenses for both BrightCloud and PAN-DB, but only one license can be active. If you're not sure whether a PAN-DB URL filtering license is active, access the firewall web interface, select **Device > Licenses**, and verify that the **Active** field displays **Yes** in the PAN-DB URL Filtering section. 3. Purchase new licenses for the firewalls that don't have valid PAN-DB licenses. Palo Alto Networks sends an email containing activation codes for the licenses you purchase. If you can't find this email, contact Customer Support before proceeding.
Step 2	Change the URL filtering vendor to PAN-DB on Panorama.	Access the Panorama web interface and perform one of the following tasks: • Change the URL Filtering Vendor on HA Panorama • Change the URL Filtering Vendor on non-HA Panorama
Step 3	Migrate the URL filtering vendor to PAN-DB on each firewall.	1. Access the firewall web interface and select **Device > Licenses**. 2. In the License Management section, select **Activate feature using authorization code**, enter the **Authorization Code**, and click **OK**. Activating the PAN-DB license automatically deactivates the BrightCloud license. 3. In the PAN-DB URL Filtering section, **Download** the seed file, select your region, and click **OK**. 4. In the Panorama web interface, click **Commit**, set the **Commit Type** to **Device Group**, select the firewall, and click **Commit** again.

Push a Policy Rule to a Subset of Firewalls

A policy *target* allows you to specify the firewalls in a device group to which to push policy rules. It allows you to exclude one or more firewalls or virtual systems, or to apply a rule only to specific firewalls or virtual systems in a device group.

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

Push a Policy Rule to a Subset of Firewalls	
Step 1 Create a rule. In this example, we define a pre-rule in the Security rulebase that permits users on the internal network to access the servers in the DMZ.	1. Select the **Policies** tab and select the **Device Group** for which you want to define a rule. 2. Select the rulebase. For this example, select **Policies > Security > Pre-Rules**. 3. Click **Add** and, in the **General** tab, enter a descriptive rule **Name**. 4. In the **Source** tab, set the **Source Zone** to **Trust**. 5. In the **Destination** tab, set the **Destination Zone** to **DMZ**. 6. In the **Service/ URL Category** tab, set the **Service** to **application-default**. 7. In the **Actions** tab, set the **Action** to **Allow**. 8. Leave all the other options at the default values.
Step 2 Target the rule to include or exclude a subset of firewalls.	To apply the rule to a selected set of firewalls: 1. Select the **Target** tab in the Policy Rule window. 2. Select the firewalls on which you want the rule to apply. If you do not select firewalls to target, the rule is added to all of the (unchecked) firewalls in the device group.  By default, although the check box for the virtual systems in the device group is unchecked, all the virtual systems will inherit the rule on commit. Select the check box for one or more virtual systems to which you want the rule to apply. 3. (Optional) To exclude a subset of firewalls from inheriting the rule, select the check box **Install on all but specified devices**.  If you select **Install on all but specified devices** and do not select any device, the rule is added to none of the firewalls in the device group. 4. Click **OK** to add the rule. 5. Save the configuration changes. a. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. b. Click **Commit**, for the **Commit Type** select **Device Group**, select the device group to which you just added the rule, and click **Commit** again.

Manage the Rule Hierarchy

The order of policy rules is critical for the security of your network. Within any policy layer (shared, device group, or locally defined rules) and rulebase (for example, shared Security pre-rules), the firewall evaluates rules from top to bottom in the order they appear in the pages of the **Policies** tab. The firewall matches a packet against the first rule that meets the defined criteria and ignores subsequent rules. Therefore, to enforce the most specific match, move the more specific rules above more generic rules.

 © Palo Alto Networks, Inc.

To understand the order in which the firewall evaluates rules by layer and by type (pre-rules, post-rules, and default rules) across the Device Group Hierarchy, see Device Group Policies.

Manage the Rule Hierarchy	
Step 1 View the rule hierarchy for each rulebase.	1. Select the **Policies** tab and click **Preview Rules**. 2. Filter the preview by **Rulebase** (for example, **Security** or **QoS**). 3. Filter the preview to display the rules of a specific **Device Group** and the rules it inherits from the Shared location and ancestor device groups. You must select a device group that has firewalls assigned to it. 4. Filter the preview by **Device** to display its locally defined rules. 5. Click the green arrow icon to apply your filter selections to the preview (see Figure: Rule Hierarchy). 6. Close the Combined Rules Preview dialog when you finish previewing rules.
Step 2 Delete or disable rules, if necessary. To determine which rules a firewall doesn't currently use, select that firewall in the **Context** drop-down on Panorama, select the rulebase (for example, **Policies > Security**), and select the **Highlight Unused Rules** check box. A dotted orange background indicates the rules that the firewall doesn't use.	1. Select the rulebase (for example, **Policies > Security > Pre Rules**) that contains the rule you will delete or disable. 2. Select the **Device Group** that contains the rule. 3. Select the rule, and click **Delete** or **Disable** as desired. Disabled rules appear in italicized font.
Step 3 Reposition rules within a rulebase, if necessary. To reposition local rules on a firewall, access its web interface by selecting that firewall in the **Context** drop-down before performing this step.	1. Select the rulebase (for example, **Policies > Security > Pre Rules**) that contains the rule you will move. 2. Select the **Device Group** that contains the rule. 3. Select the rule, select **Move**, and select: • **Move Top**—Moves the rule above all other rules in the device group (but not above rules inherited from Shared or ancestor device groups). • **Move Up**—Moves the rule above the one that precedes it (but not above rules inherited from Shared or ancestor device groups). • **Move Down**—Moves the rule below the one that follows it. • **Move Bottom**—Moves the rule below all other rules. • **Move to other device group**—See Move or Clone a Policy Rule or Object to a Different Device Group.
Step 4 If you modified the rules, save the changes.	1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 2. Click **Commit**, for the **Commit Type** select **Device Group**, select the device group that contains the rules you changed or deleted, and click **Commit** again.

Manage Templates and Template Stacks

Use templates and template stacks to define the common base configurations that enable firewalls to operate in your network. See Templates and Template Stacks for an overview of the issues you should consider when deciding which firewalls to add to which templates, ordering templates in a stack to manage layers of common and firewall group-specific settings, and overriding template settings with firewall-specific values.



To delete a template, you must first locally Disable/Remove Template Settings on the firewall. Only administrators with the superuser role can disable a template.

- ▲ Template Capabilities and Exceptions
- ▲ Add a Template
- ▲ Configure a Template Stack
- ▲ Override a Template Setting
- ▲ Disable/Remove Template Settings

Template Capabilities and Exceptions

You can use Templates and Template Stacks to define a wide array of settings, but you can perform the following tasks only locally on each managed firewall:

- Clear logs.
- Enable operational modes such as multi-vsys mode, Federal Information Processing Standards (FIPS) mode, or Common Criteria (CC) mode.
- Configure the IP addresses of a firewall HA pair.
- Configure a master key and diagnostics.
- Compare configuration files (Config Audit).



To Manage Licenses and Updates (software or content) for firewalls, use **Panorama** tab options, not templates.

Add a Template

You must add at least one template before Panorama will display the **Device** and **Network** tabs required to define the network set up and device configuration elements for firewalls.



You can avoid duplicating many configurations among templates by combining them into a template stack: see Templates and Template Stacks and Configure a Template Stack.

 © Palo Alto Networks, Inc.

Add a Template	
Step 1 Add a template.	1. Select **Panorama > Templates.** 2. Click **Add** and enter a unique **Name** to identify the template. 3. If the template has a virtual system (vsys) with configurations (for example, interfaces) that you want Panorama to push to firewalls that don't have virtual systems, select it in the **Default VSYS** drop-down. 4. In the Devices section, select check boxes to assign firewalls to the template. Whenever you add a new managed firewall to Panorama, you must assign it to the appropriate template; Panorama does not automatically assign new firewalls. When you perform a template commit, Panorama pushes the configuration to every firewall assigned to the template. 5. (Optional) Select **Group HA Peers** to display a single check box for firewalls that are in a high availability (HA) configuration. Icons indicate the HA state: green for active and yellow for passive. The firewall name of the secondary peer is in parentheses. For active/passive HA, add both peers to the same template so that both will receive the configurations. For active/active HA, whether you add both peers to the same template depends on whether each peer requires the same configurations. For a list of the configurations that PAN-OS synchronizes between HA peers, see High Availability Synchronization. 6. Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 7. Click **Commit**, for the **Commit Type** select **Template**, select the firewalls assigned to the template you just added, and click **Commit** again.
Step 2 Verify that the template is available.	After you add the first template, Panorama displays the **Device** and **Network** tabs. These tabs display a **Template** drop-down. Check that the drop-down displays the template you just added.
Step 3 Use the template to push a configuration change to firewalls.	Let's define a primary Domain Name System (DNS) server for the firewalls in the template. 1. In the **Device** tab, select the **Template** from the drop-down. 2. Select **Device > Setup > Services > Global**, and edit the Services section. 3. Enter an IP address for the **Primary DNS Server**. 4. Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 5. Click **Commit**, for the **Commit Type** select **Template**, select the firewalls assigned to the template, and click **Commit** again.

Add a Template (Continued)	
Step 4 Verify that the firewall is configured with the template settings that you pushed from Panorama.	1. In the **Context** drop-down, select one of the firewalls to which you pushed the template setting. 2. Select **Device > Setup > Services > Global**. The IP address that you pushed from the template appears. The Services section header displays a template icon (green cog) to indicate that settings in the section have values pushed from a template.

Configure a Template Stack

A template stack is a combination of templates: Panorama pushes the settings from every template in the stack to the firewalls you assign to that stack. For details and planning, see Templates and Template Stacks.

Configure a Template Stack	
Step 1 Plan the templates and their order in the stack.	For each template you will assign to the stack, Add a Template. When planning the priority order of templates within the stack (for overlapping settings), remember that Panorama doesn't check the order for invalid relationships. For example, consider a stack in which the ethernet1/1 interface is of type Layer 3 in Template_A but of type Layer 2 with a VLAN in Template_B. If Template_A has a higher priority, Panorama will push ethernet1/1 as type Layer 3 but assigned to a VLAN. Also note that a template configuration can't reference a configuration in another template, even if both templates are in the same stack. For example, a zone configuration in Template_A can't reference a zone protection profile in Template_B.
Step 2 Create a template stack.	1. Select **Panorama > Templates** and click **Add Stack**. 2. Enter a unique **Name** to identify the stack. 3. For each of the Templates the stack will combine (up to 16), click **Add** and select the template. The dialog lists the added templates in order of priority with respect to duplicate settings, where values in the higher templates override those that are lower in the list. To change the order, select a template and click **Move Up** or **Move Down**. 4. In the Devices section, select check boxes to assign firewalls. You can't assign individual virtual systems, only an entire firewall. You can assign any firewall to only one template or stack. After you finish selecting, click **OK**.

 © Palo Alto Networks, Inc.

Configure a Template Stack (Continued)		
Step 3	Edit the **Network** and **Device** settings, if necessary.  While Panorama pushes mode-specific settings only to firewalls that support those modes, this selective push doesn't adjust mode-specific values. For example, if a template has firewalls in Federal Information Processing Standards (FIPS) mode and an IKE Crypto profile that uses non-FIPS algorithms, the template commit will fail. To avoid such errors, use the **Mode** drop-down in the **Network** and **Device** tabs to filter mode-specific features and value options. In an individual firewall context, you can override settings that Panorama pushes from a stack in the same way you override settings pushed from a template: see Override a Template Setting.	1. Depending on the settings you will configure, select the **Network** or **Device** tab and select the stack in the **Template** drop-down. The tab settings are read-only when you select a stack. 2. Filter the tabs to display only the mode-specific settings you want to edit: • In the **Mode** drop-down, select or clear the **Multi VSYS**, **Operational Mode**, and **VPN Mode** filter options. • Set all the **Mode** options to reflect the mode configuration of a particular firewall by selecting it in the **Device** drop-down. 3. You can edit settings only at the template level, not at the stack level. To identify and access the template that contains the setting you want to edit: • If the page displays a table, select **Columns > Template** in the drop-down of any column header. The Template column displays the source template for each setting. If multiple templates have the same setting, the Template column displays the higher priority template. Click the template name in this column: the **Template** drop-down changes to that template, at which point you can edit the setting. • If the page doesn't display a table, hover over the template icon (green cog) for a setting: a tooltip displays the source template. If multiple templates have the same setting, the tooltip displays the higher priority template. In the **Template** drop-down, select the template that the tooltip displays to edit the setting. 4. Edit the settings as needed. 5. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 6. Click **Commit**, for the **Commit Type** select **Template**, select the firewalls assigned to the template stack, and click **Commit** again.
Step 4	Verify that the template stack works as expected.	Perform the same verification steps as when you Add a Template but select the template stack from the **Template** drop-down: 1. Use the template to push a configuration change to firewalls. 2. Verify that the firewall is configured with the template settings that you pushed from Panorama.

Override a Template Setting

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. Overrides allow for exceptions or modifications to meet your deployment needs. For example, if you

use a template to create a base configuration but a few firewalls in a test lab environment need different settings for the Domain Name System (DNS) server IP address or the Network Time Protocol (NTP) server, you can override the template settings.



If you want to disable or remove all the template or stack settings on a firewall instead of overriding a single setting, see Disable/Remove Template Settings.

Override a Template Setting		
Step 1	Access the web interface of the firewall.	Directly access the firewall by entering its IP address in the URL field of your browser, or use the **Context** drop-down in Panorama to switch to the firewall context.
Step 2	Navigate to the setting you will override. In this example, you will override the DNS server IP address that you assigned using a template in Add a Template.	1. Select **Device > Setup > Services > Global** and edit the Services section. 2. Click the template icon (green cog) for the **Primary DNS Server** to enable overrides for that field. 3. Enter a new IP address for the **Primary DNS Server**. Note that the dialog now displays a template override icon (orange-overlapping-green cogs) to indicate that the value is overridden. 4. Click **OK** and **Commit**.

Disable/Remove Template Settings

If you want to stop using a template or template stack for managing the configuration on a managed firewall, you can disable the template or stack. When disabling, you can copy the template/stack values to the local configuration of the firewall or delete the values.



If you want to override a single setting instead of disabling or removing every template or stack setting, see Override a Template Setting.

See Templates and Template Stacks for details on how to use these for managing firewalls.

Disable/Remove Template Settings	
Step 1	Access the web interface of the managed firewall as an administrator with the Superuser role. You can directly access the firewall by entering its IP address in the browser URL field or, in Panorama, select the firewall in the **Context** drop-down.
Step 2	Select **Device > Setup > Management** and edit the Panorama Settings.
Step 3	Click **Disable Device and Network Template**.
Step 4	(Optional) Select **Import Device and Network Template before disabling**, to save the configuration settings locally on the firewall. If you do not select this option, PAN-OS will delete all Panorama-pushed settings from the device.
Step 5	Click **OK** twice and then click **Commit** to save the changes.

 © Palo Alto Networks, Inc.

Transition a Firewall to Panorama Management

If you have already deployed Palo Alto Networks firewalls and configured them locally, but now want to use Panorama for centrally managing them, you must perform pre-migration planning. The migration involves importing firewall configurations into Panorama and verifying that the firewalls function as expected after the transition. If some settings are unique to individual firewalls, you can continue accessing the firewalls to manage the unique settings. You can manage any given firewall setting by pushing its value from Panorama or by configuring it locally on the firewall, but you cannot manage the setting through both Panorama and the firewall. If you want to exclude certain firewall settings from Panorama management, you can either:

- Migrate the entire firewall configuration and then, on Panorama, delete the settings that you will manage locally on firewalls. You can also Override a Template Setting that Panorama pushes to a firewall instead of deleting the setting on Panorama.
- Load a partial firewall configuration, including only the settings that you will use Panorama to manage.



Firewalls do not lose logs during the transition to Panorama management.

▲ Plan the Transition to Panorama Management

▲ Migrate a Firewall to Panorama Management

▲ Load a Partial Firewall Configuration into Panorama

Plan the Transition to Panorama Management

The following tasks are a high-level overview of the planning required to migrate firewalls to Panorama management:

❒ Decide which firewalls to migrate.

❒ Determine the Panorama and firewall software and content versions, and how you will Manage Licenses and Updates. For important details, see Panorama, Log Collector, and Firewall Version Compatibility.

❒ Plan Your Deployment for Panorama with respect to the URL filtering database (BrightCloud or PAN-DB), log collection, and administrator roles.

❒ Plan how to manage shared settings.

Plan the Device Group Hierarchy, Templates and Template Stacks in a way that will reduce redundancy and streamline the management of settings that are shared among all firewalls or within firewall sets. During the migration, you can select whether to import objects from the Shared location on the firewall into Shared on Panorama, with the following exceptions:

- If a shared firewall object has the same name and value as an existing shared Panorama object, the import excludes that firewall object.
- If the name or value of the shared firewall object differs from an existing shared Panorama object, Panorama imports the firewall object into each new device group that is created for the import.
- If a configuration imported into a template references a shared firewall object, or if a shared firewall object references a configuration imported into a template, Panorama imports the object as a shared object regardless of whether you select the **Import devices' shared objects into Panorama's shared context** check box.

- Determine if the firewall has configuration elements (policies, objects, and other settings) that you don't want to import, either because Panorama already contains similar elements or because those elements are firewall-specific (for example, timezone settings) and you won't use Panorama to manage them. You can perform a global find to determine if similar elements exist on Panorama.
- Decide the common zones for each device group. This includes a zone-naming strategy for the firewalls and virtual systems in each device group. For example, if you have zones called Branch LAN and WAN, Panorama can push policy rules that reference those zones without being aware of the variations in port or media type, platform, or logical addressing schema.
- Create a post-migration test plan.

  You will use the test plan to verify that the firewalls work as efficiently after the migration as they did before. The plan might include tasks such as:

  - Monitor the firewalls for at least 24 hours after the migration.
  - Monitor Panorama and firewall logs for anomalies.
  - Check administrator logins on Panorama.
  - Test various types of traffic from multiple sources. For example, check bandwidth graphs, session counts, and deny-rule traffic log entries (see Use Panorama for Visibility). The testing should cover a representative sample of policy configurations.
  - Check with your network operations center (NOC) and security operations center (SOC) for any user-reported issues.
  - Include any other test criteria that will help verify firewall functionality.

Migrate a Firewall to Panorama Management

When you import a firewall configuration, Panorama automatically creates a template to contain the imported network and device settings. To contain the imported policies and objects, Panorama automatically creates one device group for each firewall or one device group for each virtual system (vsys) in a multi-vsys firewall.

When you perform the following steps, Panorama imports the entire firewall configuration. Alternatively, you can Load a Partial Firewall Configuration into Panorama.



Panorama can import configurations from firewalls that run PAN-OS 5.0 or later releases and can push configurations to those firewalls. The exception is that Panorama 6.1 and later releases cannot push configurations to firewalls running PAN-OS 6.0.0 through 6.0.3.

You can import configurations from firewalls that are already managed devices but only if they are not already assigned to device groups or templates.

Migrate a Firewall to Panorama Management		
Step 1	Plan the migration.	See the checklist in Plan the Transition to Panorama Management.

 © Palo Alto Networks, Inc.

Migrate a Firewall to Panorama Management (Continued)		
Step 2	Add the firewall as a managed device.	Add a Firewall as a Managed Device: 1. Log in to Panorama, select **Panorama > Managed Devices** and click **Add**. 2. Enter the serial number of the firewall and click **OK**. If you will import multiple firewall configurations, enter the serial number of each one on a separate line. Optionally, you can copy and paste the serial numbers from a Microsoft Excel worksheet. 3. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 3	Set up a connection from the firewall to Panorama.	1. Log in to the firewall, select **Device > Setup**, and edit the Panorama Settings. 2. In the **Panorama Servers** fields, enter the IP addresses of the Panorama management server. 3. Click **OK** and **Commit**.

Migrate a Firewall to Panorama Management (Continued)	
Step 4 Import the firewall configuration into Panorama.  If you later decide to re-import a firewall configuration, first remove the firewall or its virtual systems from the device groups and template where you originally imported them. (Firewalls don't lose logs when you remove them from device groups or templates.) Because the imported policies and objects remain in the device groups, you must manually move, edit, or delete them when necessary. When re-importing, use the **Device Group Name Prefix** fields to define device group names that differ from the ones Panorama created in the original import.	1. From Panorama, select **Panorama > Setup > Operations**, click **Import device configuration to Panorama**, and select the **Device**.  Panorama can't import a configuration from a firewall that is assigned to an existing device group or template. 2. Enter a **Template Name**. If this is a multi-vsys firewall, the field is blank. Otherwise, the default value is the firewall name. You can't use the name of an existing template. 3. For a multi-vsys firewall, optionally add a character string as a **Device Group Name Prefix** for all the device groups. 4. (Optional) Edit the **Device Group** names. If this is a multi-vsys firewall, each device group has a vsys name by default. Otherwise, the default value is the firewall name. You can't use the names of existing device groups.  The **Import devices' shared objects into Panorama's shared context** check box is selected by default, which means Panorama imports objects that belong to the Shared location in the firewall to Shared in Panorama. If you clear the check box, Panorama copies shared firewall objects into device groups instead of Shared. This could create duplicate objects, so selecting the check box is a best practice in most cases. To understand the consequences of importing shared or duplicate objects into Panorama, see Plan how to manage shared settings. 5. Select a **Rule Import Location** for the imported policy rules: **Pre Rulebase** or **Post Rulebase**. Regardless of your selection, Panorama imports default security rules (intrazone-default and interzone-default) into the post-rulebase.  If Panorama has a rule with the same name as a firewall rule that you import, Panorama displays both rules. Delete one of the rules before performing a Panorama commit to prevent a commit error. 6. Click **OK**. Panorama displays the import status, result, details about your selections, details about what was imported, and any warnings. Click **Close**.
Step 5 Fine-tune the imported configuration.	1. In Panorama, select **Panorama > Config Audit**, select the **Running config** and **Candidate config** for the comparison, click **Go**, and review the output. 2. Update the device group and template configurations as needed based on the configuration audit and any warnings that Panorama displayed after the import. For example: • Delete redundant objects and policy rules. • Move or Clone a Policy Rule or Object to a Different Device Group. • Move firewalls to different device groups or templates. • Move a device group that Panorama created during the import to a different parent device group: Select **Panorama > Device Groups**, select the device group you want to move, select a new **Parent Device Group**, and click **OK**.

 © Palo Alto Networks, Inc.

Migrate a Firewall to Panorama Management (Continued)		
Step 6	Push the device configuration bundle to the firewall to remove all policy rules and objects from its local configuration. This step is necessary to prevent duplicate rule or object names, which would cause commit errors when you push the device group configuration from Panorama to the firewall in the next step.	1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. Panorama creates a device configuration bundle named <firewall_name>_import.tgz, in which all policies and objects are removed. 2. In Panorama, select **Panorama > Setup > Operations** and click **Export or push device config bundle**. 3. Select the **Device** from which you imported the configuration, click **OK**, and click **Push & Commit**. Panorama pushes the bundle and initiates a commit on the firewall.
Step 7	Push the device group and template configurations to the firewall to complete the transition to centralized management. If you are migrating multiple firewalls, perform all the preceding steps—including this one—for each firewall before continuing.	1. In Panorama, click **Commit** and for the **Commit Type** select **Device Group**. 2. Select the **Merge with Device Candidate Config**, **Include Device and Network Templates** and **Force Template Values** check boxes. 3. Select the device groups that contain the imported firewall configurations and click **Commit**.
Step 8	Consolidate all the imported firewall configurations. Required if you are migrating multiple firewalls. Settings might be duplicated among the firewalls. For example, if you imported an object with the same name from two firewalls, you must delete one object in Panorama before performing a commit on Panorama.	1. After importing all the firewall configurations, update the device groups and templates as needed to eliminate redundancy and streamline configuration management: see Fine-tune the imported configuration. (You don't need to push device configuration bundles again.) 2. Configure any firewall-specific settings. If the firewalls will have local zones, you must create them before performing a device group or template commit; Panorama can't poll the firewalls for zone name or zone configuration. If you will use local firewall rules, ensure their names are unique (not duplicated in Panorama). If necessary, you can Override a Template Setting with a firewall-specific value. 3. In Panorama, click **Commit**, for the **Commit Type** select **Device Group**, select the device groups, select the **Include Device and Network Templates** check box, and click **Commit**.
Step 9	Perform your post-migration test plan.	Perform the verification tasks that you devised during the migration planning to confirm that the firewalls work as efficiently with the Panorama-pushed configuration as they did with their original local configuration: see Create a post-migration test plan.

Load a Partial Firewall Configuration into Panorama

If some configuration settings on a firewall are common to other firewalls, you can load those specific settings into Panorama and then push them to all the other firewalls or to the firewalls in particular device groups and templates.

Load a Partial Firewall Configuration into Panorama		
Step 1	Plan the transition to Panorama.	See the checklist in Plan the Transition to Panorama Management.

Load a Partial Firewall Configuration into Panorama (Continued)		
Step 2	Resolve how to manage duplicate settings, which are those that have the same names in Panorama as in a firewall. Before you load a partial firewall configuration, Panorama and that firewall might already have duplicate settings. Loading a firewall configuration might also add settings to Panorama that are duplicates of settings in other managed firewalls. If Panorama has policy rules or objects with the same names as those on a firewall, a commit failure will occur when you try to push device group settings to that firewall. If Panorama has template settings with the same names as those on a firewall, the template values will override the firewall values when you push the template.	1. On Panorama, perform a global find to determine if duplicate settings exist. 2. Delete or rename the duplicate settings on the firewall if you will use Panorama to manage them, or delete or rename the duplicate settings on Panorama if you will use the firewall to manage them. If you will use the firewall to manage device or network settings, instead of deleting or renaming the duplicates on Panorama, you can also push the settings from Panorama (Step 6) and then Override a Template Setting on the firewall with firewall-specific values.
Step 3	Export the entire firewall configuration to your local computer.	1. On the firewall, select **Device > Setup > Operations**. 2. Click **Save named configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**. 3. Click **Export named configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. The firewall exports the configuration as an XML file.
Step 4	Import the firewall configuration snapshot into Panorama.	1. On Panorama, select **Panorama > Setup > Operations**. 2. Click **Import named Panorama configuration snapshot**, **Browse** to the firewall configuration file you exported to your computer, and click **OK**. After using this option to import a firewall configuration file, you can't use the Panorama web interface to load it. You must use the XML API or CLI, as described in the next step.

 © Palo Alto Networks, Inc.

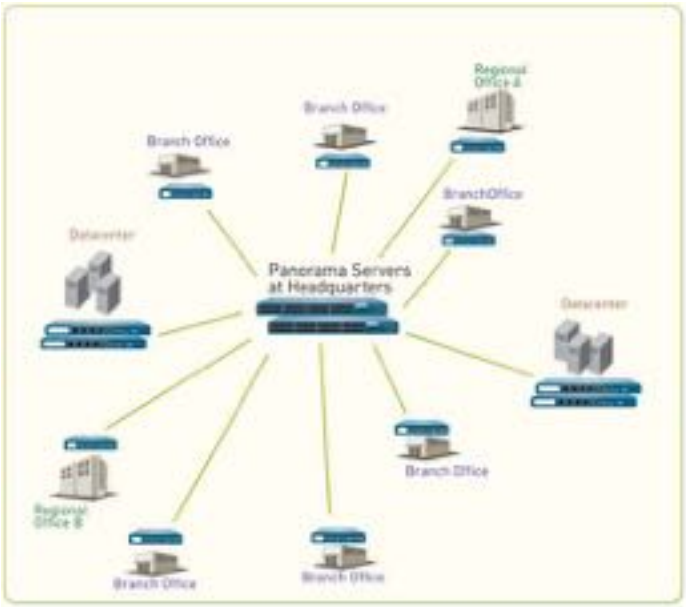
Load a Partial Firewall Configuration into Panorama (Continued)	

Step 5 Load the desired part of the firewall configuration into Panorama.
To specify a part of the configuration (for example, all application objects), you must identify the:
• Source xpath—The XML node in the firewall configuration file from which you are loading.
• Destination xpath—The node in the Panorama configuration to which you are loading.
	Use the XML API or CLI to identify and load the partial configuration:
	1. Use the firewall XML API or CLI to identify the source xpath.
	For example, the xpath for application objects in vsys1 of the firewall is:
	`/config/devices/entry[@name='localhost.localdomain']/vsys/entry[@name='vsys1']/application`
	2. Use the Panorama XML API or CLI to identify the destination xpath.
	For example, to load application objects into a device group named US-West, the xpath is:
	`/config/devices/entry[@name='localhost.localdomain']/device-group/entry[@name='US-West']/application`
	3. Use the Panorama CLI to load the configuration and commit the change:
	# **load config partial from** <filename> **from-xpath** <source-xpath> **to-xpath** <destination-xpath> **mode** [append|merge|replace]
	For example, enter the following to load the application objects from vsys1 on an imported firewall configuration named fw1-config.xml into a device group named US-West on Panorama:
	# load config partial from fw1-config.xml from-xpath devices/entry[@name='localhost.localdomain']/vsys/entry[@name='vsys1']/application to-xpath /config/devices/entry[@name='localhost.localdomain']/device-group/entry[@name='US-West']/application mode merge
	# commit

Step 6 Push the partial configuration from Panorama to the firewall to complete the transition to centralized management.	1. On the firewall, delete any rules or objects that have the same names as those in Panorama. If the device group for that firewall has other firewalls with rules or objects that are duplicated in Panorama, perform this step on those firewalls also. For details, see Step 2. 2. On Panorama, click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 3. On Panorama, click **Commit** and for the **Commit Type** select **Device Group**. 4. Select the **Merge with Device Candidate Config**, **Include Device and Network Templates** and **Force Template Values** check boxes. 5. Select the device groups that contain the imported firewall configurations and click **Commit**. 6. If the firewall has a device or network setting that you won't use Panorama to manage, Override a Template Setting on the firewall.
Step 7 Perform your post-migration test plan.	Perform the verification tasks that you devised during the migration planning to confirm that the firewall works as efficiently with the Panorama-pushed configuration as it did with its original local configuration: see Create a post-migration test plan.

Use Case: Configure Firewalls Using Panorama

Let's say that you want to use Panorama in a high availability configuration to manage a dozen firewalls on your network: you have six firewalls deployed across six branch offices, a pair of firewalls in a high availability configuration at each of two data centers, and a firewall in each of the two regional head offices.



The first step in creating your central management strategy is to determine how to group the firewalls into device groups and templates to efficiently push configurations from Panorama. You can base the grouping on the business functions, geographic locations, or administrative domains of the firewalls. In this example, you create two device groups and three templates to administer the devices using Panorama:

- ▲ Device Groups
- ▲ Templates
- ▲ Set Up Your Centralized Configuration and Policies

Device Groups

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

- DG_BranchAndRegional for grouping devices that serve as the security gateways at the branch offices and at the regional head offices. We placed the branch office firewalls and the regional office firewalls in the same device group because devices with similar functions will require similar policy rulebases.
- DG_DataCenter for grouping the devices that secure the servers at the data centers.

We can then administer shared policy rules across both device groups as well as administer distinct device group rules for the regional office and branch office groups. Then for added flexibility, the local administrator at a regional or branch office can create local rules that match specific source, destination, and service flows for accessing applications and services that are required for that office. In this example, we create the following hierarchy for security rules. you can use a similar approach for any of the other rulebases.

 © Palo Alto Networks, Inc.

Device Groups	DG_BranchAndRegional		DG_DataCenter
Rules	**Regional**	**Branch**	**Datacenter**
Shared pre-rule	Allow DNS and SNMP services.		
	Acceptable use policy that denies access to specified URL categories and peer-to-peer traffic that is of risk level 3, 4, and 5.		
Device Group pre-rule	Allow Facebook to all users in the marketing group in the regional offices only.		Allow access to the Amazon cloud application for the specified hosts/servers in the datacenter.
Local rules on a device	None		
Device Group post-rule	None		
Shared post-rule	To enable logging for all Internet-bound traffic on your network, create a rule that allows or denies all traffic from the trust zone to the untrust zone.		

Templates

When grouping firewalls for templates, we must take into account the differences in the networking configuration. For example, if the interface configuration is not the same—the interfaces are unlike in type, or the interfaces used are not alike in the numbering scheme and link capacity, or the zone to interface mappings are different —the firewalls must be in separate templates. Further, the way the firewalls are configured to access network resources might be different because the firewalls are spread geographically; for example, the DNS server, syslog servers and gateways that they access might be different. So, to allow for an optimal base configuration, you must place the firewalls in separate templates as follows:

- T_Branch for the branch office firewalls
- T_Regional for the regional office firewalls
- T_DataCenter for the data center firewalls



If you plan to deploy your firewalls in an active/active HA configuration, assign each firewall in the HA pair to a separate template. Doing so gives you the flexibility to set up separate networking configurations for each peer. For example, you can manage the networking configurations in a separate template for each peer so that each can connect to different northbound and southbound routers, and can have different OSPF or BGP peering configurations.

Set Up Your Centralized Configuration and Policies

Using the example described in the preceding topics (starting with Use Case: Configure Firewalls Using Panorama), perform the following tasks to centrally deploy and administer firewalls:

- TASK 1—Add the firewalls as managed devices, and deploy content updates and PAN-OS software updates to those firewalls.
- TASK 2—Use templates to administer a base configuration.
- TASK 3—Use device groups to manage the firewall policies.
- TASK 4—Preview your rules and commit your changes to Panorama, device groups, and templates.

Add the Managed Firewalls and Deploy Updates		
TASK 1 Add the firewalls as managed devices, and deploy content updates and PAN-OS software updates to those firewalls.		
Step 1	For each firewall that Panorama will manage, Add a Firewall as a Managed Device.	In this example, add 12 firewalls.
Step 2	Deploy the content updates to the firewalls. If you purchased a Threat Prevention subscription, the content and antivirus databases are available to you. First install the **Applications** or **Applications and Threats** database, then the **Antivirus**. To review the status or progress for all tasks performed on Panorama, see View Panorama Task Completion History.	1. Select **Panorama > Device Deployment > Dynamic Updates**. 2. Click **Check Now** to check for the latest updates. If the value in the Action column is **Download**, this indicates an update is available. 3. Click **Download**. When the download completes, the value in the Action column changes to **Install**. 4. In the **Action** column, click **Install**. Use the filters or user-defined tags to select the managed firewalls on which you would like to install this update. 5. Click **OK**, then monitor the status, progress, and result of the content update for each firewall. The **Result** column displays the success or failure of the installation.

 © Palo Alto Networks, Inc.

Add the Managed Firewalls and Deploy Updates (Continued)		
Step 3	Deploy the software updates to the firewalls.	1. Select **Panorama > Device Deployment > Software**. 2. Click **Check Now** to check for the latest updates. If the value in the Action column is **Download**, this indicates an update is available. 3. Locate the version that you need for each hardware model and click **Download**. When the download completes, the value in the Action column changes to **Install**. 4. In the Action column, click the **Install** link. Use the filters or user-defined tags to select the managed firewalls on which to install this version. 5. Enable the check box for **Reboot device after install** or **Upload only to device (do not install)** and click **OK**. The **Results** column displays the success or failure of the installation.

Use Templates to Administer a Base configuration		
TASK 2 Use templates to administer a base configuration.		
Step 1	For each template you will use, Add a Template and assign the appropriate firewalls to each.	In this example, create templates named T_Branch, T_Regional, and T_DataCenter.
Step 2	Define a DNS server, NTP server, syslog server, and login banner. Repeat this step for each template.	1. In the **Device** tab, select the **Template** from the drop-down. 2. Define the DNS and NTP servers: a. Select **Device > Setup > Services > Global** and edit the Services. b. In the **Services** tab, enter an IP address for the **Primary DNS Server**. For any firewall that has more than one virtual system (vsys), for each vsys, add a DNS server profile to the template (**Device > Server Profiles > DNS**). c. In the **NTP** tab, enter an IP address for the **Primary NTP Server**. d. Click **OK** to save your changes. 3. Add a login banner: select **Device > Setup > Management**, edit the General Settings, enter text for the **Login Banner** and click **OK**. 4. Configure a Syslog server profile (**Device > Server Profiles > Syslog**).
Step 3	Enable HTTPS, SSH, and SNMP access to the management interface of the managed firewalls. Repeat this step for each template.	1. In the **Device** tab, select the **Template** from the drop-down. 2. Select **Setup > Management**, and edit the Management Interface Settings. 3. Under Services, select the **HTTPS**, **SSH**, and **SNMP** check boxes, and click **OK**.

Use Templates to Administer a Base configuration (Continued)	
Step 4 Create a Zone Protection profile for the firewalls in the data center template (T_DataCenter).	1. Select the **Network** tab and, in the **Template** drop-down, select T_DataCenter. 2. Select **Network Profiles > Zone Protection** and click **Add**. 3. For this example, enable protection against a SYN flood—In the **Flood Protection** tab, select the **SYN** check box, set the **Action** to **SYN Cookies** as, set the **Alert** packets/second to `100`, set the **Activate** packets/second to `1000`, and set the **Maximum** packets/second to `10000`. 4. For this example, enable alerts—In the **Reconnaissance Protection** tab, select the **Enable** check boxes for **TCP Port Scan**, **Host Sweep**, and **UDP Port Scan**. Ensure the Action values are set to **alert** (the default value). 5. Click **OK** to save the Zone Protection profile.
Step 5 Configure the interface and zone settings in the data center template (T_DataCenter), and then attach the Zone Protection profile you just created. Before performing this step, you must have configured the interfaces locally on the firewalls. As a minimum, for each interface, you must have defined the interface type, assigned it to a virtual router (if needed), and attached a security zone.	1. Select the **Network** tab and, in the **Template** drop-down, select T_DataCenter. 2. Select **Network > Interface** and, in the Interface column, click the interface name. 3. Select the **Interface Type** from the drop-down. 4. In the **Virtual Router** drop-down, click **New Virtual Router**. When defining the router, ensure the **Name** matches what is defined on the firewall. 5. In the **Security Zone** drop-down, click **New Zone**. When defining the zone, ensure that the **Name** matches what is defined on the firewall. 6. Click **OK** to save your changes to the interface. 7. Select **Network > Zones**, and select the zone you just created. Verify that the correct interface is attached to the zone. 8. In the **Zone Protection Profile** drop-down, select the profile you created, and click **OK**.
Step 6 Commit your template changes.	1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 2. Click **Commit**, for the **Commit Type** select **Template**, select the firewalls assigned to the templates in which you made changes, and click **Commit** again.

Use Device Groups to Push Policy Rules	
TASK 3 Use device groups to manage the policy rules on your firewalls.	
Step 1 Create device groups and assign the appropriate firewalls to each device group: see Add a Device Group.	In this example, create device groups named DG_BranchAndRegional and DG_DataCenter. When configuring the DG_BranchAndRegional device group, you must assign a **Master** firewall. This is the only firewall in the device group that gathers user and group mapping information for policy evaluation.

 © Palo Alto Networks, Inc.

Use Device Groups to Push Policy Rules (Continued)		
Step 2	Create a shared pre-rule to allow DNS and SNMP services.	1. Create a shared application group for the DNS and SNMP services. a. Select **Objects > Application Group** and click **Add**. b. Enter a **Name** and select the **Shared** check box to create a shared application group object. c. Click **Add**, type `DNS`, and select **dns** from the list. Repeat for SNMP and select **snmp**, **snmp-trap**. d. Click **OK** to create the application group. 2. Create the shared rule. a. Select the **Policies** tab and, in the **Device Group** drop-down, select **Shared**. b. Select the **Security > Pre-Rules** rulebase. c. Click **Add** and enter a **Name** for the security rule. d. In the **Source** and **Destination** tabs for the rule, click **Add** and enter a **Source Zone** and a **Destination Zone** for the traffic. e. In the **Applications** tab, click **Add**, type the name of the applications group object you just created, and select it from the drop-down. f. In the **Actions** tab, set the **Action** to **Allow**, and click **OK**.
Step 3	Define the corporate acceptable use policy for all offices. In this example, create a shared rule that restricts access to some URL categories and denies access to peer-to-peer traffic that is of risk level 3, 4, or 5.	1. Select the **Policies** tab and, in the **Device Group** drop-down, select **Shared**. 2. Select **Security > Pre-Rules** and click **Add**. 3. In the **General** tab, enter a **Name** for the security rule. 4. In the **Source** and **Destination** tabs, click **Add** and select **any** for the traffic **Source Zone** and **Destination Zone**. 5. In the **Application** tab, define the application filter: a. Click **Add** and click **New Application Filter** in the footer of the drop-down. b. Enter a **Name**, and select the **Shared** check box. c. In the Risk column, select levels **3**, **4**, and **5**. d. In the Technology column, select **peer-to-peer**. e. Click **OK** to save the new filter. 6. In the **Service/URL Category** tab, URL Category section, click **Add** and select the categories you want to block (for example, **streaming-media**, **dating**, and **online-personal-storage**). 7. You can also attach the default URL Filtering profile—In the **Actions** tab, Profile Setting section, select the **Profile Type** option **Profiles**, and select the **URL Filtering** option **default**. 8. Click **OK** to save the security pre-rule.

Use Device Groups to Push Policy Rules (Continued)		
Step 4	Allow Facebook for all users in the Marketing group in the regional offices only. Enabling a security rule based on user and group has the following prerequisite tasks: • Set up User-ID on the firewalls. • Enable User-ID for each zone that contains the users you want to identify. • Define a master firewall for the DG_BranchAndRegional device group (Step 1).	1. Select the **Policies** tab and, in the **Device Group** drop-down, select DG_BranchAndRegional. 2. Select the **Security > Pre-Rules** rulebase. 3. Click **Add** and enter a **Name** for the security rule. 4. In the **Source** tab, **Add** the Source Zone that contains the Marketing group users. 5. In the **Destination** tab, **Add** the Destination Zone. 6. In the **User** tab, **Add** the Marketing user group to the Source User list. 7. In the **Application** tab, click **Add**, type `Facebook`, and then select it from the drop-down. 8. In the **Action** tab, set the **Action** to **Allow**. 9. In the **Target** tab, select the regional office firewalls and click **OK**.
Step 5	Allow access to the Amazon cloud application for the specified hosts/servers in the data center.	1. Create an address object for the servers/hosts in the data center that need access to the Amazon cloud application. a. Select **Objects > Addresses** and, in the **Device Group** drop-down, select DG_DataCenter. b. Click **Add** and enter a **Name** for the address object. c. Select the **Type**, and specify an IP address and netmask (**IP Netmask**), range of IP addresses (**IP Range**), or **FQDN**. d. Click **OK** to save the object. 2. Create a security rule that allows access to the Amazon cloud application. a. Select **Policies > Security > Pre-Rules** and, in the **Device Group** drop-down, select DG_DataCenter. b. Click **Add** and enter a **Name** for the security rule. c. Select the **Source** tab, **Add** the Source Zone for the data center, and **Add** the address object (Source Address) you just defined. d. Select the **Destination** tab and **Add** the Destination Zone. e. Select the **Application** tab, click **Add**, type `amazon`, and select the Amazon applications from the list. f. Select the **Action** tab and set the **Action** to **Allow**. g. Click **OK** to save the rule.
Step 6	To enable logging for all Internet-bound traffic on your network, create a rule that matches trust zone to untrust zone.	1. Select the **Policies** tab and, in the **Device Group** drop-down, select **Shared**. 2. Select the **Security > Pre-Rules** rulebase. 3. Click **Add** and enter a **Name** for the security rule. 4. In the **Source** and **Destination** tabs for the rule, **Add** `trust_zone` as the Source Zone and `untrust_zone` as the Destination Zone. 5. In the **Action** tab, set the **Action** to **Deny**, set the **Log Setting** to **Log at Session end**, and click **OK**.

 © Palo Alto Networks, Inc.

Preview the Rules and Commit Changes	
TASK 4	
Preview your rules and commit your changes to Panorama, device groups, and templates.	
Step 1	In the **Policies** tab, click **Preview Rules**, and select a **Rulebase**, **Device Group**, and **Device**. This preview enables you to visually evaluate how rules are layered for a particular rulebase. Close the preview dialog when you are done.
Step 2	Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 3	Click **Commit**, for the **Commit Type** select **Device Group**, select the device groups you added, select the **Include Device and Network Templates** check box, and click **Commit** again.
Step 4	In the **Context** drop-down, select the firewall to access its web interface and confirm that Panorama applied the template and policy configurations.

 © Palo Alto Networks, Inc.

Manage Log Collection

All Palo Alto Networks next-generation firewalls can generate logs that provide an audit trail of firewall activities. For Centralized Logging and Reporting, you must forward the logs generated on the firewalls to Panorama. You can then configure Panorama to aggregate the logs and forward them to remote logging destinations. If you forward logs to a Panorama virtual appliance, you don't need to perform any additional tasks to enable logging. If you will forward logs to an M-Series appliance in Panorama mode or Log Collector mode, you must add the Log Collectors as managed collectors and assign them to Collector Groups to access, manage, and update the Log Collectors using Panorama. To determine which deployment best suits your needs, see Plan a Log Collection Deployment.



To manage the System and Config logs that Panorama generates locally, see Monitor Panorama.

- Configure a Managed Collector
- Manage Collector Groups
- Configure Log Forwarding to Panorama
- Verify Log Forwarding to Panorama
- Modify Log Forwarding and Buffering Defaults.
- Configure Log Forwarding from Panorama to External Destinations
- Log Collection Deployments

Configure a Managed Collector

To enable the Panorama management server (Panorama virtual appliance or M-Series appliance in Panorama mode) to manage a Log Collector, you must add it as a managed collector. The M-Series appliance in Panorama mode has a predefined (default) local Log Collector. However, switching from Panorama Mode to Log Collector Mode would remove the local Log Collector and would require you to re-configure the appliance as a Dedicated Log Collector (M-Series appliance in Log Collector mode). When the Panorama management server has a high availability (HA) configuration, each HA peer can have a local Log Collector. Dedicated Log Collectors don't support HA.

We recommend that you install the same Applications update on Panorama as on managed Collectors and firewalls. For details, see Panorama, Log Collector, and Firewall Version Compatibility.

We recommend retaining a local Log Collector and local Collector Group on the M-Series appliance in Panorama mode, regardless of whether it manages Dedicated Log Collectors.

Configure a Managed Collector	
Step 1 Perform initial setup of the M-Series appliance in Log Collector mode if you haven't already. Only Dedicated Log Collectors require this step.	1. Rack mount the M-Series appliance. Refer to the M-100 or M-500 Hardware Reference Guide for instructions. 2. Perform Initial Configuration of the M-Series Appliance. When configuring interfaces, configure only the Management (MGT) interface. Switching to Log Collector mode (in the next step) removes any Eth1 and Eth2 interface configurations. If the Log Collector will use Eth1 and Eth2, add them when configuring the Log Collector later in this procedure. 3. Register Panorama and Install Licenses. 4. Install Content and Software Updates for Panorama. 5. Configure each logging disk pair. This task is required to make the RAID disks available for logging. Optionally, you can add disks to Increase Storage on the M-Series Appliance.

Step 2 Switch from Panorama mode to Log Collector mode.
Switching to Log Collector mode reboots the appliance, deletes any existing log data, and deletes all configurations except management access settings. Switching the mode does not delete licenses, software updates, or content updates.
	1. Access the CLI of the M-Series appliance.
	2. Enter the following command to switch to Log Collector mode:
	`> request system system-mode Logger`
	3. Enter **Y** to confirm the mode change. The M-Series appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the M-Series appliance to display the Panorama login prompt.
	If you see a **CMS Login** prompt, press Enter without typing a username or password.
	4. Log back in to the CLI.
	5. Enter the following command to verify that the switch to Log Collector mode succeeded:
	`> show system info | match system-mode`
	If the mode change succeeded, the output displays:
	`system-mode: logger`

 © Palo Alto Networks, Inc.

Configure a Managed Collector (Continued)		
Step 3	Enable connectivity among the M-Series appliances.	These steps vary by Log Collector type. For HA deployments, <IPaddress1> and <IPaddress2> are for the management interface of the primary and secondary Panorama management server respectively. For non-HA deployments, specify only <IPaddress1>. • Dedicated Log Collectors—Run the following commands at the CLI of each Log Collector: **> Configure** **# set deviceconfig system panorama-server <IPaddress1> panorama-server-2 <IPaddress2>** **# commit** • Local Log Collectors—These steps are required only for an HA deployment: a. Log into the CLI of the primary Panorama and enter: **> Configure** **# set deviceconfig system panorama-server <IPaddress2>** **# commit** b. Log into the CLI of the secondary Panorama and enter: **> Configure** **# set deviceconfig system panorama-server <IPaddress1>** **# commit**
Step 4	Record the serial number of the Log Collector. You will need this when you add the Log Collector as a managed collector.	The steps to display the serial number vary by Log Collector type: • Local—Access the Panorama web interface and record the value on the **Dashboard** tab, General Information section, **Serial #** field. In an HA deployment, record the **Serial #** of each Panorama peer on which you will configure a Log Collector. • Dedicated—Access the Log Collector CLI, run the **show system info** command, and record the serial number.
Step 5	Configure the general settings of the Log Collector.	Use the web interface of the primary Panorama management server to perform these steps: 1. Select **Panorama > Managed Collectors** and **Add** a new Log Collector or edit the predefined local Log Collector (named **default**). Although the secondary Panorama HA peer has a predefined local Log Collector, you must manually add it on the primary Panorama. 2. In the **General** tab, **Collector S/N** field, enter the serial number you recorded for the Log Collector.

Configure a Managed Collector (Continued)		
Step 6	Configure network access for the Log Collector. Perform this step only for a Dedicated Log Collector or a local Log Collector on the secondary Panorama HA peer. Although you defined similar parameters during initial configuration of the Panorama management server, you must re-define the parameters for the Log Collector.	1. In the **Panorama Server IP** field, enter the IP address or FQDN of the solitary (non-HA) or primary (HA) Panorama. For an HA deployment, enter the IP address or FQDN of the secondary Panorama peer in the **Panorama Server IP 2** field. These fields are required. 2. Configure the IP addresses of the **Primary DNS Server** and **Secondary DNS Server**. 3. (Optional) Set the **Timezone** that Panorama will use to record log entries.
Step 7	Configure administrative access to the Log Collector CLI. Only Dedicated Log Collectors require this step. The default CLI administrator is **admin**. You cannot modify this username nor add CLI administrators.	1. Select the **Authentication** tab, select the password **Mode**, and enter the **Password** (the default is `admin`). 2. Enter the number of **Failed Attempts** to log in that Panorama allows before locking out the administrator. Enter the **Lockout Time** in minutes. These settings can help protect the Log Collector from a brute force attack.
Step 8	Configure the Log Collector interfaces. Perform this step only for a Dedicated Log Collector or a local Log Collector on the secondary Panorama HA peer. The **Eth1** or **Eth2** interfaces are available only if you defined them during initial configuration of the Panorama management server.	1. Configure each interface that the Log Collector will use. Only the **Management** interface is required. For each interface, select the corresponding tab and configure one or both of the following field sets based on the IP protocols of your network. • IPv4—**IP Address**, **Netmask**, and **Default Gateway** • IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway** 2. (Optional) In the **Management** tab, select the **SNMP** service if you will use SNMP to monitor the Log Collector. Using SNMP requires additional steps besides configuring the Log Collector. For details, see Monitor Panorama and Log Collector Statistics Using SNMP. 3. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again. This step is required before you can enable logging disks or assign the Eth1 and Eth2 interfaces to logging functions. 4. (Optional) Edit the Log Collector and select the interfaces (**mgmt**, **eth1**, or **eth2**) that it will use for **Device Log Collection** and **Collector Group Communication** (default is **mgmt**).
Step 9	Enable the logging disks.	1. Select **Disks** and **Add** each disk pair. 2. Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

 © Palo Alto Networks, Inc.

Configure a Managed Collector (Continued)	
Step 10 Verify your changes.	1. Verify that the **Panorama > Managed Collectors** page lists the Log Collector you added. The Connected column displays a check mark to indicate that the Log Collector is connected to Panorama. You might have to wait a few minutes before the page displays the updated connection status. Until you Configure a Collector Group and perform a Collector Group commit, the Configuration Status column displays Out of Sync, the Run Time Status column displays disconnected, and the CLI command `show interface all` displays the interfaces as `down`. 2. Click **Statistics** in the last column to verify that the logging disks are enabled.
Step 11 Next steps...	Before a Log Collector can receive firewall logs, you must: 1. Configure Log Forwarding to Panorama. 2. Configure a Collector Group. The predefined local Log Collector is preassigned to a predefined Collector Group.

Manage Collector Groups

A Collector Group is one or more Log Collectors that operate as a single logical unit for collecting firewall logs. You can configure a Collector Group with multiple Log Collectors to ensure log redundancy or to accommodate logging rates that exceed the capacity of a single Log Collector (see Panorama Platforms). To understand the risks and recommended mitigations, see Caveats for a Collector Group with Multiple Log Collectors.

The M-Series appliance in Panorama mode (Panorama management server) has a predefined (default) local Collector Group that contains a predefined local Log Collector. However, switching to Log Collector mode would remove the local Log Collector and Collector Group; you would have to Set up the M-Series Appliance as a Log Collector, add it as a managed collector to the Panorama management server, and configure a Collector Group to contain the managed collector.



If you delete a Collector Group, you will lose logs.
We recommend retaining a local Log Collector and local Collector Group on the M-Series appliance in Panorama mode, regardless of whether it manages Dedicated Log Collectors.

- Configure a Collector Group
- Move a Log Collector to a Different Collector Group
- Remove a Firewall from a Collector Group

Configure a Collector Group

After you configure Log Collectors and firewalls, you must assign them to a Collector Group so that the firewalls can send logs to the Log Collectors. A Collector Group with multiple Log Collectors has the following requirements:

- All the Log Collectors in any particular Collector Group must be the same platform: all M-100 appliances or all M-500 appliances.
- Log redundancy is available only if each Log Collector has the same number of logging disks. To add disks to a Log Collector, see Increase Storage on the M-Series Appliance.

 © Palo Alto Networks, Inc.

Configure a Collector Group	
Step 1 Perform the following tasks before configuring the Collector Group.	In these tasks, skip any steps that involve configuring or committing changes to the Collector Group; you will perform those steps later in the current procedure. 1. Add a Firewall as a Managed Device for each firewall that you will assign to the Collector Group. 2. (Optional) Configure Log Forwarding from Panorama to External Destinations. 3. Configure a Managed Collector for each Log Collector that you will assign to the Collector Group. You must manually add each Dedicated Log Collector (M-Series appliance in Log Collector mode). The M-Series appliance in Panorama mode has a predefined local Log Collector that you don't need to add. If you will use SNMP for monitoring, select the **SNMP** service when you configure the **Management** interface of a Log Collector. Using SNMP requires additional steps besides configuring the Collector Group. For details, see Monitor Panorama and Log Collector Statistics Using SNMP.
Step 2 Add the Collector Group.	1. Access the Panorama web interface, select **Panorama > Collector Groups**, and **Add** a Collector Group or edit an existing one. The M-Series appliance in Panorama mode has a predefined Collector Group named **default**. 2. In the **General** tab, enter a **Name** for the Collector Group if you are adding one. You cannot rename an existing Collector Group. 3. Enter the **Minimum Retention Period** in days (1-2,000) for which the Collector Group will retain firewall logs. 4. (Optional) **Enable log redundancy across collectors** to ensure that no logs are lost if any one Log Collector in the Collector Group becomes unavailable. Each log will have two copies and each copy will reside on a different Log Collector. Redundancy is available only if each Log Collector has the same number of logging disks. Enabling redundancy creates more logs and therefore requires more storage capacity. When a Collector Group runs out of space, it deletes older logs. Redundancy also doubles the log processing traffic in a Collector Group, which reduces its maximum logging rate by half, as each Log Collector must distribute a copy of each log it receives. If you add multiple Log Collectors to a single Collector group, enabling redundancy is a best practice.

Configure a Collector Group (Continued)	
Step 3 (Optional) Configure SNMP monitoring.	Select **Monitoring**, select the SNMP **Version** and enter the corresponding details: • **V2c**—Enter the **SNMP Community String**, which identifies a community of SNMP managers and monitored devices (Log Collectors, in this case), and serves as a password to authenticate the community members to each other. Don't use the default community string `public`; it is well known and therefore not secure. • **V3**—Create at least one SNMP view group and one user. User accounts and views provide authentication, privacy, and access control when Log Collectors forward traps and SNMP managers get Log Collector statistics. • **Views**—Each view is a paired OID and bitwise mask: the OID specifies a MIB and the mask (in hexadecimal format) specifies which objects are accessible within (include matching) or outside (exclude matching) that MIB. Click **Add** in the first list and enter a **Name** for the group of views. For each view in the group, click **Add** and configure the view **Name**, **OID**, matching **Option** (**include** or **exclude**), and **Mask**. • **Users**—Click **Add** in the second list, enter a username in the Users column, select the **View** group from the drop-down, enter the authentication password (**Auth Password**) used to authenticate to the SNMP manager, and enter the privacy password (**Priv Password**) used to encrypt SNMP messages to the SNMP manager.
Step 4 Assign Log Collectors and firewalls to the Collector Group.	1. Select the **Device Log Forwarding** tab. 2. In the Collector Group Members section, **Add** the Log Collectors. All the Log Collectors for any particular Collector Group must be the same platform (all M-100 appliances or all M-500 appliances). 3. In the Log Forwarding Preferences section, click **Add**. 4. In the Devices section, click **Modify**, select the firewalls, and click **OK**. You cannot assign PA-7050 firewalls to a Collector Group. However, when you monitor logs or generate reports for a device group that includes a PA-7050 firewall, Panorama queries the firewall in real-time to display its log data. 5. In the Collectors section, **Add** the Log Collectors to which the firewalls will forward logs. If you assign multiple Log Collectors, the first one will be the primary; if the primary becomes unavailable, the firewalls send logs to the next Log Collector in the list. To change the priority of a Log Collector, select it and **Move Up** (higher priority) or **Move Down** (lower priority). 6. Click **OK**.

 © Palo Alto Networks, Inc.

Configure a Collector Group (Continued)	
Step 5 Define the storage capacity (log quotas) and expiration period for each log type.	1. Return to the **General** tab and click the **Log Storage** value. If the field displays 0MB, verify that you enabled the disk pairs for logging and committed the changes (see Configure a Managed Collector, **Disks** tab). 2. Enter the log storage **Quota(%)** for each log type. 3. Enter the **Max Days** (expiration period) for each log type (range is 1-2,000). By default, the fields are blank for all log types, which means the logs never expire.
Step 6 (Optional) Configure log forwarding from the Collector Group to external services. To perform this step, you must have added server profiles for the external services in the task Configure Log Forwarding from Panorama to External Destinations. In a high availability (HA) deployment, you can configure each Panorama HA peer to forward logs to different external services. For details, see Deploy Panorama with Default Log Collectors.	1. Select the **Collector Log Forwarding** tab. 2. For each log Severity level in the **System**, **Threat**, and **Correlation** tabs, click a cell in the SNMP Trap, Email Profile, or Syslog Profile column, and select the server profile. 3. In the **Config**, **HIP Match**, and **Traffic** tabs, select the **SNMP Trap**, **Email**, or **Syslog** server profile. 4. For each Verdict in the **WildFire** tab, click a cell in the SNMP Trap, Email Profile, or Syslog Profile column, and select the server profile. 5. Click **OK** to save the Collector Group.
Step 7 Commit the changes and verify that the Log Collectors you assigned to the Collector Group are connected to, and synchronized with, Panorama.	1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 2. Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Group you added, and click **Commit** again. 3. Select **Panorama > Managed Collectors**. The Connected column displays a check mark icon to indicate that a Log Collector is connected to Panorama. The Configuration Status column indicates whether the configurations you committed to Panorama and the Log Collectors are synchronized (green icon) or are not synchronized (red icon) with each other. The Collector Group won't receive firewall logs until you Configure Log Forwarding to Panorama.

Move a Log Collector to a Different Collector Group

When you Plan a Log Collection Deployment, you assign Log Collectors to a Collector Group based on the logging rate and log storage requirements of that Collector Group. If the rates and required storage increase in a Collector Group, the best practice is to Increase Storage on the M-Series Appliance or Configure a Collector Group with additional Log Collectors. However, in some deployments, it might be more economical to move Log Collectors between Collector Groups.

The log data on a Log Collector becomes inaccessible after you remove it from a Collector Group. Also, you must perform a factory reset on the Log Collector before adding it to another Collector Group; a factory reset removes all configuration settings and logs.

When a Log Collector is local to an M-Series appliance in Panorama mode, move it only if the M-Series appliance is the passive peer in a high availability (HA) configuration. HA synchronization will restore the configurations that the factory reset removes. Never move a Log Collector when it's local to an M-Series appliance that is the active HA peer.

All the Log Collectors in any particular Collector Group must be the same platform: all M-100 appliances or all M-500 appliances.

Move a Log Collector to Different Collector Group	
Step 1 Remove the Log Collector from Panorama management.	1. Select **Panorama > Collector Groups** and select the Collector Group that contains the Log Collector you will move. 2. Select the **Device Log Forwarding** tab and, in the Log Forwarding Preferences list, perform the following steps for each set of firewalls assigned to the Log Collector you will move: a. In the Devices column, click the link for the firewalls assigned to the Log Collector. b. In the Collectors column, select the Log Collector and click **Delete**. To reassign the firewalls, **Add** the new Log Collector to which they will forward logs. c. Click **OK** twice to save your changes. 3. Select **Panorama > Managed Collectors**, select the Log Collector you will move, and click **Delete**. 4. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 5. Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Group from which you deleted the Log Collector, and click **Commit** again.
Step 2 Reset the Log Collector to its factory default settings. Do not interrupt the factory reset or reboot processes. Otherwise, you might render the M-Series appliance unusable.	1. Log in to the CLI of the Log Collector. 2. Enter the following CLI operational command: `> debug system maintenance-mode` The Log Collector takes approximately six minutes to reboot in maintenance mode. 3. After the Log Collector reboots, press `Enter` to access the maintenance mode menu. 4. Select **`Factory Reset`** and press Enter. 5. Select **`Factory Reset`** and press Enter again. The factory reset and subsequent reboot take approximately eight minutes in total, after which the Log Collector won't have any configuration settings or log data. The default username and password to log in to the Log Collector is admin/admin.

 © Palo Alto Networks, Inc.

Move a Log Collector to Different Collector Group (Continued)		
Step 3	Reconfigure the Log Collector.	1. Perform Initial Configuration of the M-Series Appliance. 2. Register Panorama and Install Licenses. 3. Install Content and Software Updates for Panorama. 4. Set up the M-Series Appliance as a Log Collector.
Step 4	Configure a Collector Group.	Add the Log Collector to its new Collector Group and assign firewalls to the Log Collector. When you commit the Collector Group configuration, Panorama starts redistributing logs across the Log Collectors. This process can take hours for each terabyte of logs. During the redistribution process, the maximum logging rate is reduced. In the **Panorama > Collector Groups** page, the Redistribution State column indicates the completion status of the process as a percentage.

Remove a Firewall from a Collector Group

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

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

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

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

Configure Log Forwarding to Panorama

By default, firewalls store all log files locally. To aggregate logs on Panorama, you must configure the firewalls to forward logs to Panorama. Before starting this procedure, you must Add a Device Group and Add a Template for the firewalls that will forward logs.



To forward firewall logs directly to external services (for example, a syslog server) and also to Panorama, see Configure Log Forwarding.

For details about all the log collection deployments that Panorama supports, see Log Forwarding Options.

The PA-7000 Series firewall can't forward logs to Panorama, only to external services. However, when you monitor logs or generate reports for a device group that includes a PA-7000 Series firewall, Panorama queries the firewall in real-time to display its log data.

If Panorama will manage firewalls running software versions earlier than PAN-OS 7.0, specify a WildFire server from which Panorama can gather analysis information for WildFire samples that those firewalls submit. Panorama uses the information to complete WildFire Submissions logs that are missing field values introduced in PAN-OS 7.0. Firewalls running earlier releases won't populate those fields. To specify the server, select **Panorama > Setup > WildFire**, edit the General Settings, and enter the **WildFire Server** name. The default is **wildfire-public-cloud**, which is the WildFire cloud hosted in the United States.

Configure Log Forwarding to Panorama		
Step 1	Create a log forwarding profile. The profile defines the destination of Traffic, Threat, and WildFire logs. (Threat logs include URL Filtering and Data Filtering logs.)	1. Select **Objects > Log Forwarding** and select the **Device Group** of the firewalls that will forward logs. 2. Click **Add** and enter a **Name** to identify the profile. 3. For each log type and each severity level or WildFire verdict, select the **Panorama** check box. 4. Click **OK** to save the profile.
Step 2	Assign the log forwarding profile to security rules. To trigger log generation and forwarding, the rules require certain security profiles according to log type: • Traffic logs—No security profile is necessary; the traffic only needs to match a specific security rule. • Threat logs—The traffic must match any security profile assigned to a security rule. • WildFire logs—The traffic must match a WildFire Analysis profile assigned to a security rule.	Perform the following steps for each rule that will trigger log forwarding: 1. Select the rulebase of the rule that will trigger log forwarding (for example, **Policies > Security > Pre Rules**), select the **Device Group** of the firewalls that will forward logs, and select the rule. 2. Select the **Actions** tab and select the **Log Forwarding** profile you just created. 3. In the **Profile Type** drop-down, select **Profiles** or **Group**, and then select the security profiles or **Group Profile** required to trigger log generation and forwarding. 4. For Traffic logs, select one or both of the **Log At Session Start** and **Log At Session End** check boxes, and click **OK**.

 © Palo Alto Networks, Inc.

Configure Log Forwarding to Panorama (Continued)		
Step 3	Configure the destination of System, Config, and HIP Match logs.  You cannot forward Correlation logs (correlated events) from the firewalls to Panorama. On the logs that are forwarded from your managed firewalls, Panorama matches for the conditions specified in the correlation objects and automatically generates correlated event(s) when a match is observed. If you want, you can then forward these correlated events (Correlation logs) from Panorama to an external syslog server.	1. Select **Device > Log Settings** and select the **Template** of the firewalls that will forward logs. 2. For System logs, click each Severity level, select the **Panorama** check box, and click **OK**. 3. For Config and HIP Match logs, click the edit icon, select the **Panorama** check box, and click **OK**.
Step 4	(M-Series appliances only) Configure Panorama to receive the logs.	1. For each Log Collector that will receive logs, Configure a Managed Collector. 2. Configure a Collector Group, in which you assign firewalls to specific Log Collectors for log forwarding.
Step 5	Commit your configuration changes.	1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 2. Click **Commit**, for the **Commit Type** select **Device Group**, select the device group of the firewalls that will forward logs, select the **Include Device and Network Templates** check box, and click **Commit** again. 3. Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Group you just configured to receive the logs, and click **Commit** again. 4. Verify Log Forwarding to Panorama to confirm that your configuration is successful.  To change the log forwarding mode that the firewalls use to send logs to Panorama and to specify which Panorama HA peer can receive logs, you can Modify Log Forwarding and Buffering Defaults. You can also Manage Storage Quotas and Expiration Periods for Logs and Reports.

Verify Log Forwarding to Panorama

After you Configure Log Forwarding to Panorama, test that your configuration succeeded.

Verify Log Forwarding to Panorama	
Step 1	Access the firewall CLI.

Step 2	If you configured Log Collectors, verify that each firewall has a log forwarding preference list.
	`> show log-collector preference-list`
	If the Collector Group has only one Log Collector, the output will look something like this:
	`Log collector Preference List`
	`Serial Number: 003001000024`
	`IP Address:10.2.133.48`

Step 3	Verify that each firewall is forwarding logs. `> show logging-status device <firewall-serial-number>` For successful forwarding, the output indicates that the log forwarding agent is active. For a Panorama virtual appliance, the agent is “Panorama.” For an M-Series appliance, the agent is a “Log Collector.”
Step 4	View the average logging rate. The displayed rate will be the average logs/second for the last five minutes. • If Log Collectors receive the logs, access the Panorama web interface, select **Panorama > Managed Collectors** and click the **Statistics** link in the far-right column. • If a Panorama virtual appliance receives the logs, access the Panorama CLI and run the following command: `debug log-collector log-collection-stats show incoming-logs` This command also works on an M-Series appliance.

 © Palo Alto Networks, Inc.

Modify Log Forwarding and Buffering Defaults

You can define the log forwarding mode that the firewalls use to send logs to Panorama and, when configured in a high availability (HA) configuration, specify which Panorama peer can receive logs. To access these options, select **Panorama > Setup > Management**, edit the Logging and Reporting Settings, and select the **Log Export and Reporting** tab.

- Define the log forwarding mode on the firewall: The firewalls can forward logs to Panorama (pertains to both the M-Series appliance and the Panorama virtual appliance) in either Buffered Log Forwarding mode or in the Live Mode Log Forwarding mode.

Logging Options	Description
Buffered Log Forwarding from Device Default: Enabled  It is a best practice to select the **Buffered Log Forwarding from Device** option.	Allows each managed firewall to buffer logs and send the logs at 30-second intervals to Panorama (not user configurable). Buffered log forwarding is very valuable when the firewall loses connectivity to Panorama. The firewall buffers log entries to its local hard disk and keeps a pointer to record the last log entry that was sent to Panorama. When connectivity is restored the firewall resumes forwarding logs from where it left off. The disk space available for buffering depends on the log storage quota for the platform and the volume of logs that are pending roll over. If the firewall was disconnected for a long time and the last log forwarded was rolled over, all the logs from its local hard disk will be forwarded to Panorama on reconnection. If the available space on the local hard disk of the firewall is consumed, the oldest entries are deleted to allow logging of new events.
Live Mode Log Forwarding from Device This option is enabled when the check box for **Buffered Log Forwarding from Device** is cleared.	In live mode, the managed firewall sends every log transaction to Panorama at the same time as it records it on the firewall.

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

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

Logging Options	Pertains to	Description
Get Only New Logs on Convert to Primary Default: Disabled	Panorama virtual appliance that is mounted to a Network File System (NFS) datastore and is set up in a high availability (HA) configuration	With NFS logging, when you have a pair of Panorama servers configured in a high availability configuration, only the primary Panorama peer mounts the NFS datastore. Therefore, the firewalls can only send logs to the primary Panorama peer, which can write to the NFS datastore. When an HA failover occurs, the **Get Only New Logs on Convert to Primary** option allows an administrator to configure the managed firewalls to send only newly generated logs to Panorama. This event is triggered when the priority of the active-secondary Panorama is promoted to primary and it can begin logging to the NFS. This behavior is typically enabled to prevent the firewalls from sending a large volume of buffered logs when connectivity to Panorama is restored after a significant period of time.

 © Palo Alto Networks, Inc.

Configure Log Forwarding from Panorama to External Destinations

Panorama enables you to forward logs to external servers, including syslog, email, and SNMP trap servers. Forwarding firewall logs from Panorama reduces the load on the firewalls and provides a reliable and streamlined approach to forwarding logs to remote destinations. You can also forward logs that Panorama and its managed collectors generate.



To forward firewall logs directly to external services and also to Panorama, see Configure Log Forwarding.

For details about all the log collection deployments that Panorama supports, see Log Forwarding Options.

On a Panorama virtual appliance running Panorama 5.1 or earlier releases, you can use Secure Copy (SCP) commands from the CLI to export the entire log database to an SCP server and import it to another Panorama virtual appliance. A Panorama virtual appliance running Panorama 6.0 or later releases, and M-Series appliances running any release, do not support these options because the log database on those platforms is too large for an export or import to be practical.

Configure Log Forwarding from Panorama to External Destinations		
Step 1	Configure the firewalls to forward logs to Panorama.	Configure Log Forwarding to Panorama.
Step 2	Configure a server profile for each external service that will receive log data.	1. Select **Panorama > Server Profiles** and select the type of server that will receive the log data: **SNMP Trap**, **Syslog**, or **Email**. 2. Configure the server profile. Optionally, you can configure separate profiles for different log types and severity levels or WildFire verdicts. • Configure an SNMP Trap server profile. For details on how Simple Network Management Protocol (SNMP) works for Panorama and Log Collectors, refer to SNMP for Palo Alto Networks Devices. • Configure a Syslog server profile. If the syslog server requires client authentication, use the **Panorama > Certificate Management > Certificates** page to create a certificate for securing syslog communication over SSL. • Configure an Email server profile.
Step 3	Configure the destinations for: • Firewall logs that a Panorama virtual appliance collects. • Logs that Panorama (a virtual appliance or M-Series appliance) and managed collectors generate.	1. Select **Panorama > Log Settings**. 2. For System, Correlation, and Threat logs, click each Severity level, select the **SNMP Trap**, **Email**, or **Syslog** server profile you just created, and click **OK**. 3. For WildFire logs, click each Verdict, select the **SNMP Trap**, **Email**, or **Syslog** server profile you just created, and click **OK**. 4. For Config, HIP Match, and Traffic logs, click the edit icon, select the **SNMP Trap**, **Email**, or **Syslog** server profile you just created, and click **OK**.

© Palo Alto Networks, Inc.

Configure Log Forwarding from Panorama to External Destinations (Continued)		
Step 4	(M-Series appliance only) Configure the destinations for firewall logs that an M-Series appliance in Panorama or Log Collector mode collects.  Each Collector Group can forward logs to different destinations. If the Log Collectors are local to a high availability (HA) pair of M-Series appliances in Panorama mode, you must log into each HA peer to configure log forwarding for its Collector Group.	1. Select **Panorama > Collector Groups** and select the Collector Group that receives the firewall logs. 2. Select the **Collector Log Forwarding** tab. 3. For each log Severity level in the **System**, **Threat**, and **Correlation** tabs, click a cell in the SNMP Trap, Email Profile, or Syslog Profile column, and select the server profile you just created. 4. In the **Config**, **HIP Match**, and **Traffic** tabs, select the **SNMP Trap**, **Email**, or **Syslog** server profile you just created. 5. For each Verdict in the **WildFire** tab, click a cell in the SNMP Trap, Email Profile, or Syslog Profile column, and select the server profile you just created. 6. Click **OK** to save your changes to the Collector Group.
Step 5	(SNMP trap forwarding only) Enable your SNMP manager to interpret traps.	Load the Supported MIBs for Palo Alto Networks devices and, if necessary, compile them. For the specific steps, refer to the documentation of your SNMP manager.
Step 6	(Syslog forwarding only) If the syslog server requires client authentication, and the firewalls forward logs to M-Series appliances in Log Collector mode, assign a certificate that secures syslog communication over SSL.	Perform the following steps for each M-Series appliance in Log Collector mode: 1. Select **Panorama > Managed Collectors** and select the Log Collector. 2. In the **General** tab, select the **Certificate for Secure Syslog**, and click **OK**.
Step 7	Commit your configuration changes.	1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 2. Click **Commit**, for the **Commit Type** select **Device Group**, select all the device groups of the firewalls from which Panorama collects logs, select the **Include Device and Network Templates** check box, and click **Commit** again. 3. (M-Series appliance only) Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Group you just configured to forward logs, and click **Commit** again.
Step 8	(Optional) Verify the external services are receiving logs from Panorama.	• Email server—Verify that the specified recipients are receiving logs as email notifications. • Syslog server—Refer to the documentation for your syslog server to verify it is receiving logs as syslog messages. • SNMP manager—Use an SNMP Manager to Explore MIBs and Objects to verify it is receiving logs as SNMP traps.

 © Palo Alto Networks, Inc.

Log Collection Deployments

The following topics describe how to configure log collection in the most typical deployments.



The deployments in these topics all describe Panorama in a high availability (HA) configuration. Palo Alto Networks recommends HA because it enables automatic recovery (in case of server failure) of components that are not saved as part of configuration backups. In HA deployments, the Panorama management server only supports an active/passive configuration.

- Plan a Log Collection Deployment
- Deploy Panorama with Dedicated Log Collectors
- Deploy Panorama with Default Log Collectors
- Deploy Panorama Virtual Appliances with Local Log Collection

Plan a Log Collection Deployment

- Panorama and Log Collector Platforms
- Collector Groups with Single or Multiple Log Collectors
- Log Forwarding Options

Panorama and Log Collector Platforms

Decide which Panorama Platforms to use for the Panorama management server and Log Collectors based on the logging rate and geographic distribution of managed firewalls.



If you initially implement log collection using the default Log Collectors but later require more storage or higher logging rates than these support, you can switch to a deployment with Dedicated Log Collectors (M-Series appliances in Log Collector mode). You can also implement a hybrid deployment that includes both default and Dedicated Log Collectors. However, if you initially implement log collection using Dedicated Log Collectors, you will lose logs if you later switch to a deployment that involves only the default Log Collectors because of the reduced storage capacity.

If you deploy firewalls remotely, consider the bandwidth requirement for the connection between the firewalls and Panorama, in addition to whether Panorama supports the required logging rate. Deploying Dedicated Log Collectors close to the firewalls can increase the bandwidth for log forwarding.

The following table summarizes your choice of Log Collector when considering the rate at which it receives firewall logs.

Logging Rate	Log Collector
Up to 10,000 logs/second	Depends on the Panorama management server: • Virtual appliance—Panorama collects logs without any Log Collector. • M-Series appliance—Local default Log Collector

Logging Rate	Log Collector
Up to 30,000 logs/second	M-100 appliance in Log Collector Mode. Each M-100 appliance can process up to 30,000 logs/second and store up to 4TB of log data.
Up to 60,000 logs/second	M-500 appliance in Log Collector Mode. Each M-500 appliance can process up to 60,000 logs/second and store up to 8TB of log data.

Collector Groups with Single or Multiple Log Collectors

You can configure a Collector Group with multiple Log Collectors to ensure log redundancy, increase the log retention period, or accommodate logging rates that exceed the capacity of a single Log Collector (see Panorama Platforms for capacity information). To understand the requirements, risks and recommended mitigations, see Caveats for a Collector Group with Multiple Log Collectors.

All the Log Collectors in any particular Collector Group must be the same platform: all M-500 appliances or all M-100 appliances. Log redundancy is available only if each Log Collector has the same number of logging disks.

Log Forwarding Options

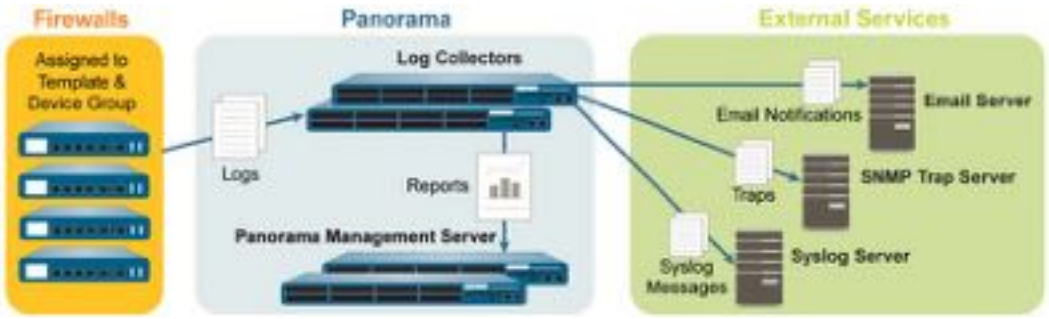
By default, each firewall stores its log files locally. To use Panorama for centralized log monitoring and report generation, you must Configure Log Forwarding to Panorama. You can also Configure Log Forwarding from Panorama to External Destinations for archiving, notification, or analysis. When forwarding from Panorama, you can include the System and Config logs that Panorama and its Log Collectors generate. External services include syslog servers, email servers, or SNMP trap servers. The firewall, Panorama virtual appliance, or M-Series appliance that forwards the logs to external services converts the logs to the appropriate format (syslog message, email notification, or SNMP trap).

Palo Alto Networks devices support the following log forwarding options:

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

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

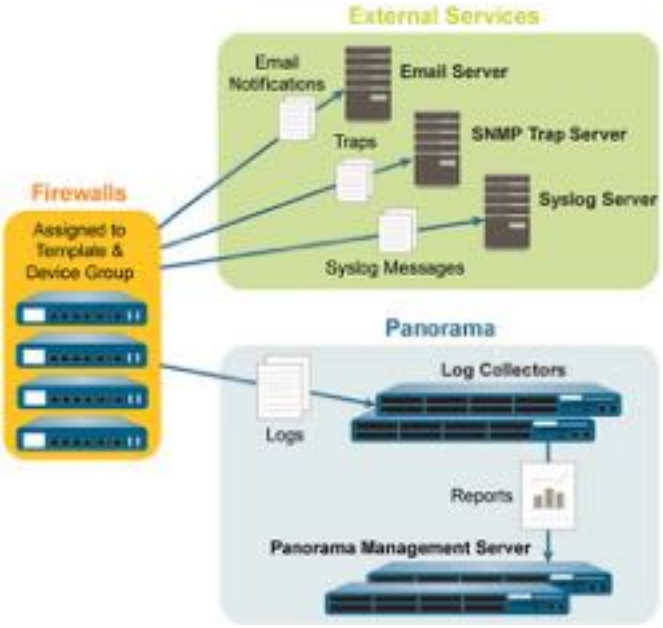
Figure: Log Forwarding to Panorama and then to External Services



 © Palo Alto Networks, Inc.

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

Figure: Log Forwarding to External Services and Panorama in Parallel



- Forward logs from firewalls directly to external services and also from Panorama to external services—This configuration is a hybrid of the previous two and is best for deployments that require sending syslog messages to multiple Security Information and Event Management (SIEM) solutions, each with its own message format (for example, Splunk and ArcSight). This duplicate forwarding doesn't apply to SNMP traps or email notifications.

Deploy Panorama with Dedicated Log Collectors

The following figures illustrate Panorama in a Distributed Log Collection Deployment. In these examples, the Panorama management server comprises two M-Series appliances in Panorama mode that are deployed in an active/passive high availability (HA) configuration. Alternatively, you can use an HA pair of Panorama virtual appliances. The firewalls send logs to Dedicated Log Collectors (M-Series appliances in Log Collector mode). This is the recommended configuration if the firewalls generate over 10,000 logs/second. (For details on deployment options, see Plan a Log Collection Deployment.)



If you will assign more than one Log Collector to a Collector Group, see Caveats for a Collector Group with Multiple Log Collectors to understand the requirements, risks, and recommended mitigations.

Figure: Single Dedicated Log Collector Per Collector Group

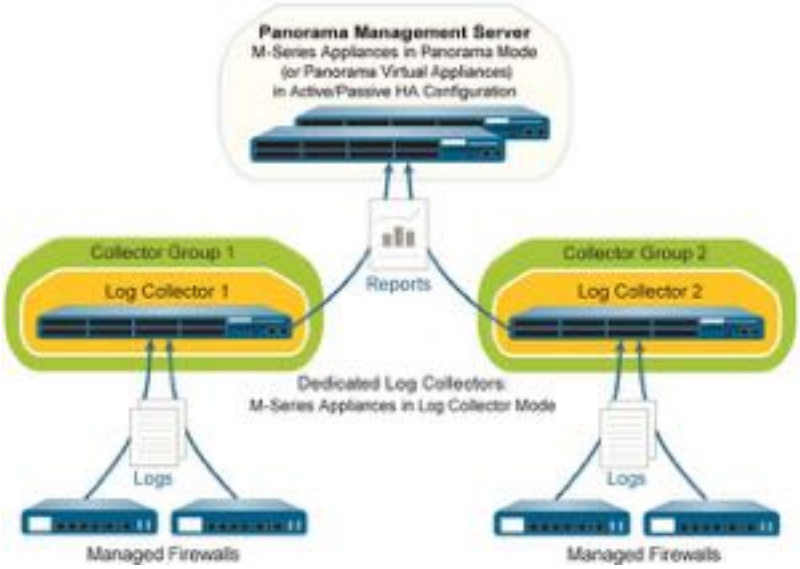
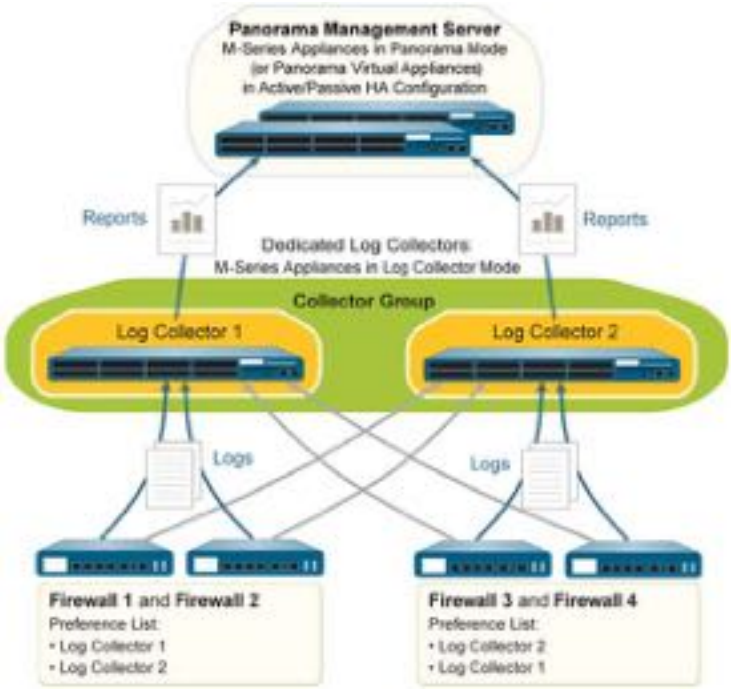


Figure: Multiple Dedicated Log Collectors Per Collector Group



Perform the following steps to deploy Panorama with Dedicated Log Collectors. Skip any steps you have already performed (for example, the initial setup).

 © Palo Alto Networks, Inc.

Deploy Panorama with Dedicated Log Collectors	
Step 1 Perform the initial setup of the Panorama management server (virtual appliances or M-Series appliances) and the Dedicated Log Collectors.	For each M-Series appliance: 1. Rack mount the M-Series appliance. Refer to the M-100 or M-500 Hardware Reference Guide for instructions. 2. Perform Initial Configuration of the M-Series Appliance.  To reduce the traffic load on the management (MGT) interface and to improve security for management traffic, Palo Alto Networks recommends configuring Panorama to use the Eth1 and Eth2 interfaces for log collection and Collector Group communication. You define these interfaces during initial configuration of the Panorama management server. 3. Configure each logging disk pair. This task is required to make the RAID disks available for logging. Optionally, you can add disks to Increase Storage on the M-Series Appliance. 4. Register Panorama and Install Licenses. 5. Install Content and Software Updates for Panorama. For each virtual appliance (if any): 1. Install the Panorama Virtual Appliance. 2. Perform Initial Configuration of the Panorama Virtual Appliance. 3. Register Panorama and Install Licenses. 4. Install Content and Software Updates for Panorama. For the Panorama management server (virtual appliance or M-Series appliance), you must also Set Up HA on Panorama.

Deploy Panorama with Dedicated Log Collectors (Continued)	

Step 2 Switch from Panorama mode to Log Collector mode on each M-Series appliance that will be a Dedicated Log Collector.
Switching the mode of an M-Series appliance deletes any existing log data and deletes all configurations except the management access settings. After the switch, the M-Series appliance retains CLI access but loses web interface access.
	1. Connect to the M-Series appliance in one of the following ways:
	• Attach a serial cable from your computer to the Console port on the M-Series appliance. Then use terminal emulation software (9600-8-N-1) to connect.
	• Use terminal emulation software such as PuTTY to open an SSH session to the IP address that you specified for the MGT interface of the M-Series appliance during initial configuration.
	2. Log in to the CLI when prompted. Use the default admin account and the password that you specified during initial configuration.
	3. Switch to Log Collector mode by entering the following command:
	`> request system system-mode logger`
	4. Enter **Y** to confirm the mode change. The M-Series appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the M-Series appliance to see the Panorama login prompt.
	If you see a **CMS Login** prompt, this means the Log Collector has not finished rebooting. Press Enter at the prompt without typing a username or password.
	5. Log back in to the CLI.
	6. Verify that the switch to Log Collector mode succeeded:
	`> show system info | match system-mode`
	If the mode change succeeded, the output displays:
	`system-mode: logger`

Step 3 Record the serial number of each Log Collector. You need the serial numbers to add the Log Collectors as managed collectors on the Panorama management server.	1. At the CLI of each Log Collector, enter the following command to display its serial number. `> show system info | match serial` 2. Record the serial number.

 © Palo Alto Networks, Inc.

Deploy Panorama with Dedicated Log Collectors (Continued)		
Step 4	Add each Log Collector as a managed collector.	Use the web interface of the primary Panorama management server peer to Configure a Managed Collector: 1. Select **Panorama > Managed Collectors**, click **Add**, and enter the serial number you recorded for the Log Collector in the **General** tab, **Collector S/N** field. 2. Enter the IP address or FQDN of the primary and secondary Panorama HA peers in the **Panorama Server IP** field and **Panorama Server IP 2** field respectively. These fields are required. 3. Select the **Management** tab and complete one or both of the following field sets for the management interface, depending on the IP protocols of your network: • IPv4—**IP Address**, **Netmask**, and **Default Gateway** • IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway** 4. (Optional) Select the **SNMP** check box if you will use SNMP to monitor the Log Collector. Using SNMP requires additional steps besides configuring the Log Collector. For details, see Monitor Panorama and Log Collector Statistics Using SNMP. 5. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again. This step is required before you can enable logging disks on the Log Collectors. 6. Verify that the **Panorama > Managed Collectors** page lists the Log Collector you added. The Connected column displays a check mark to indicate that the Log Collector is connected to Panorama. You might have to wait a few minutes before the page displays the updated connection status. At this point, the Configuration Status column displays Out of Sync and the Run Time Status column displays disconnected. The status will change to In Sync and connected after you configure a Collector Group.

Step 5	Enable connectivity between each Log Collector and the Panorama management server.
	This step is required before you can enable logging disks on the Log Collectors.
		Enter the following commands at the CLI of each Log Collector. `<IPaddress1>` is for the MGT interface of the active Panorama and `<IPaddress2>` is for the MGT interface of the passive Panorama.
		`> configure`
		`# set deviceconfig system panorama-server <IPaddress1> panorama-server-2 <IPaddress2>`
		`# commit`
		`# exit`

Step 6	Enable the logging disks on each Log Collector.	Use the web interface of the primary Panorama management server peer to perform these steps: 1. Select **Panorama > Managed Collectors** and edit the Log Collector. 2. Select **Disks** and **Add** each disk pair. 3. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again.

Deploy Panorama with Dedicated Log Collectors (Continued)	
Step 7 (Optional) Configure the **Eth1** and/or **Eth2** interfaces if the Log Collectors will use them for log collection and Collector Group communication. These interfaces are available only if you configured them for the Panorama management server during initial configuration.  Palo Alto Networks recommends using **Eth1** and/or **Eth2** to reduce the traffic load on the MGT interface and to improve security for management traffic.	Use the web interface of the primary Panorama management server peer to perform these steps for each Log Collector: 1. Select **Panorama > Managed Collectors** and edit the Log Collector. 2. Configure the network settings of the **Eth1** and/or **Eth2** interfaces. For each interface, select the corresponding tab and configure one or both of the following field sets based on the IP protocols of your network. • IPv4—**IP Address**, **Netmask**, and **Default Gateway** • IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway** 3. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again. This step is required before you can assign the Eth1 and Eth2 interfaces to logging functions. 4. Select **Panorama > Managed Collectors** and edit the Log Collector. 5. Select the interfaces (MGT, Eth1, or Eth2) that the Log Collector will use for **Device Log Collection** and **Collector Group Communication** (default is MGT). 6. Click **OK** to save your changes.
Step 8 Add a Firewall as a Managed Device.	Use the web interface of the primary Panorama management server peer to perform this task for each firewall that will forward logs to Log Collectors.

 © Palo Alto Networks, Inc.

Deploy Panorama with Dedicated Log Collectors (Continued)	
Step 9 Configure the Collector Group. If each Collector Group will have one Log Collector, repeat this step for each Collector Group before continuing. If you will assign all the Log Collectors to one Collector Group, perform this step only once.	Use the web interface of the primary Panorama management server to Configure a Collector Group: 1. Select **Panorama > Collector Groups**, click **Add**, and enter a **Name** for the Collector Group. If you add multiple Log Collectors to a single Collector group, selecting the **Enable log redundancy across collectors** check box is a best practice. For redundancy, each Log Collector requires the same number of disks. 2. (Optional) Select the **Monitoring** tab and configure the settings if you will use SNMP to monitor Log Collectors. 3. Select the **Device Log Forwarding** tab and, in the Collector Group Members section, assign one or more Log Collectors. All the Log Collectors for any particular Collector Group must be the same platform: all M-100 appliances or all M-500 appliances. 4. In the Log Forwarding Preferences section, assign firewalls according to the number of Log Collectors in this Collector Group: • Single—Assign the firewalls that will forward logs to that Log Collector, as illustrated in Figure: Single Dedicated Log Collector Per Collector Group. • Multiple—Assign each firewall to both Log Collectors for redundancy. When you configure the preferences, make Log Collector 1 the first priority for half the firewalls and make Log Collector 2 the first priority for the other half, as illustrated in Figure: Multiple Dedicated Log Collectors Per Collector Group. 5. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again. 6. Click **Commit**, set the **Commit Type** to **Collector Group**, select the Collector Groups you added, and click **Commit** again. 7. Select **Panorama > Managed Collectors** to verify that the Log Collector configuration is synchronized with Panorama. The Configuration Status column should display In Sync and the Run Time Status column should display connected.
Step 10 Configure log forwarding.	Use the web interface of the primary Panorama management server peer to: 1. Configure Log Forwarding to Panorama. 2. Verify Log Forwarding to Panorama. 3. (Optional) Configure Log Forwarding from Panorama to External Destinations.

Deploy Panorama with Default Log Collectors

The following figures illustrate Panorama in a centralized log collection deployment. In these examples, the Panorama management server comprises two M-Series appliances in Panorama mode that are deployed in an active/passive high availability (HA) configuration. The firewalls send logs to the predefined (default) local Log Collector on each Panorama M-Series appliance. This is the recommended deployment if the firewalls generate up to 10,000 logs/second. (For details on deployment options, see Plan a Log Collection Deployment.)

If you will assign more than one Log Collector to a Collector Group, see Caveats for a Collector Group with Multiple Log Collectors to understand the requirements, risks, and recommended mitigations.

After implementing this deployment, if the logging rate increases beyond 10,000 logs per second, it is recommended that you add Dedicated Log Collectors (M-Series appliances in Log Collector mode) as described in Deploy Panorama with Dedicated Log Collectors. Such an expansion might require reassigning firewalls from the default Log Collectors to Dedicated Log Collectors.

Figure: Single Default Log Collector Per Collector Group

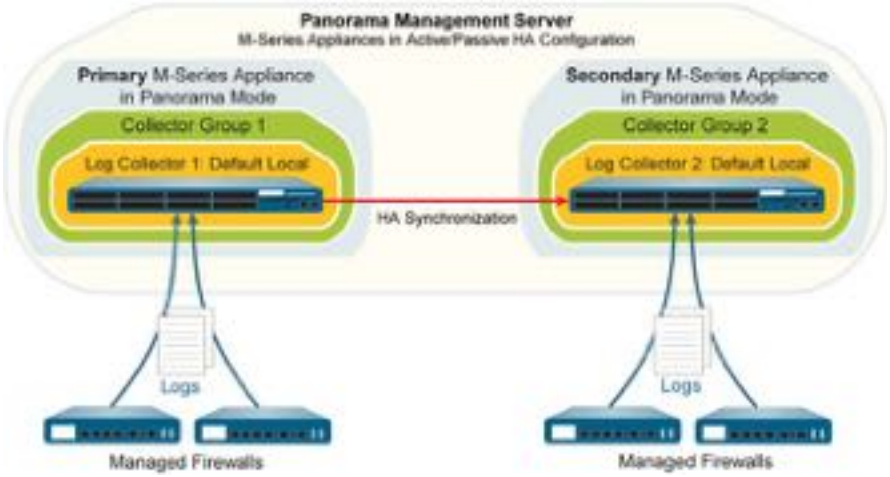
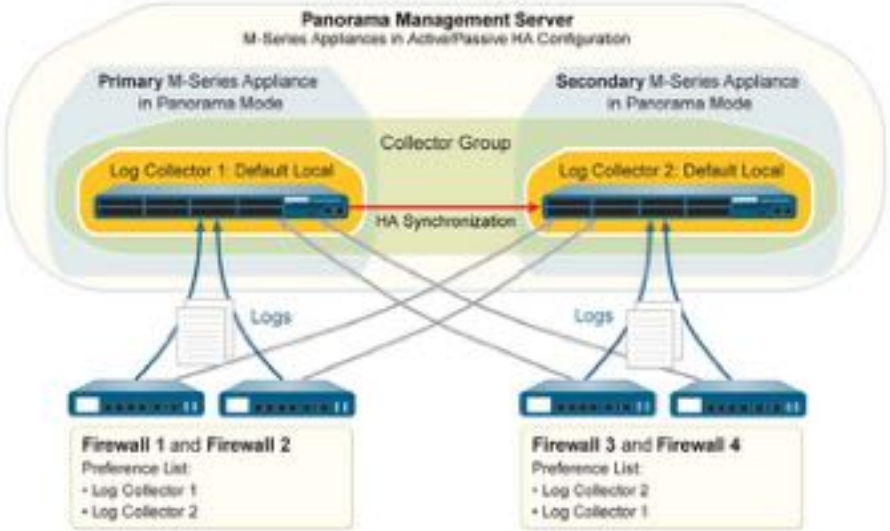


Figure: Multiple Default Log Collectors Per Collector Group



 © Palo Alto Networks, Inc.

Perform the following steps to deploy Panorama with default Log Collectors. Skip any steps you have already performed (for example, the initial setup).

Deploy Panorama with Default Log Collectors	
Step 1 Perform the initial setup of each M-Series appliance.	1. Rack mount the M-Series appliance. Refer to the M-100 or M-500 Hardware Reference Guide for instructions. 2. Perform Initial Configuration of the M-Series Appliance.  To reduce the traffic load on the management (MGT) interface and to improve security for management traffic, Palo Alto Networks recommends configuring Panorama to use the Eth1 and Eth2 interfaces for log collection and Collector Group communication. 3. Configure each logging disk pair. This task is required to make the RAID disks available for logging. Optionally, you can add disks to Increase Storage on the M-Series Appliance. 4. Register Panorama and Install Licenses. 5. Install Content and Software Updates for Panorama. 6. Set Up HA on Panorama.

Deploy Panorama with Default Log Collectors (Continued)	

Step 2 Perform the following steps to prepare Panorama for log collection.	1. Connect to the primary Panorama in one of the following ways:
	• Attach a serial cable from your computer to the Console port on the primary Panorama. Then use terminal emulation software (9600-8-N-1) to connect.
	• Use terminal emulation software such as PuTTY to open an SSH session to the IP address that you specified for the MGT interface of the primary Panorama during initial configuration.
	2. Log in to the CLI when prompted. Use the default admin account and the password that you specified during initial configuration.
	3. Enable the primary Panorama to connect to the secondary Panorama by entering the following command, where `<IPaddress2>` represents the MGT interface of the secondary Panorama:
	`> configure`
	`# set deviceconfig system panorama-server <IPaddress2>`
	`# commit`
	4. Log in to the CLI of the secondary Panorama.
	5. Enable the secondary Panorama to connect to the primary Panorama by entering the following command, where `<IPaddress1>` represents the MGT interface of the primary Panorama:
	`> configure`
	`# set deviceconfig system panorama-server <IPaddress1>`
	`# commit`
	`# exit`
	6. In the CLI of the secondary Panorama, enter the following command to display the serial number, and then record it:
	`> show system info | match serial`
	You need the serial number to add the Log Collector of the secondary Panorama as a managed collector to the primary Panorama.

Step 3 Edit the Log Collector that is local to the primary Panorama.	Use the web interface of the primary Panorama to perform these steps: 1. Select **Panorama > Managed Collectors** and select the **default** Log Collector. 2. (Optional) Select the interfaces (MGT, Eth1, or Eth2) to use for **Device Log Collection** and **Collector Group Communication**. By default, Panorama uses the MGT interface for these functions. The Eth1 and Eth2 interfaces are available only if you configured them for the Panorama management server during initial configuration. 3. Select **Disks** and **Add** each logging disk pair. 4. Click **OK** to save your changes.

 © Palo Alto Networks, Inc.

Deploy Panorama with Default Log Collectors (Continued)	
Step 4 Configure the Log Collector that is local to the secondary Panorama.  Panorama treats this Log Collector as remote because it's not local to the primary Panorama. Therefore you must manually add it on the primary Panorama.	Use the web interface of the primary Panorama to Configure a Managed Collector: 1. Select **Panorama > Managed Collectors** and **Add** the Log Collector. 2. Enter the serial number (**Collector S/N**) you recorded for the Log Collector of the secondary Panorama. 3. Enter the IP address or FQDN of the primary and secondary Panorama HA peers in the **Panorama Server IP** field and **Panorama Server IP 2** field respectively. These fields are required. 4. Configure the network settings of each interface that the Log Collector will use. The **Management** interface is required. If you configured the Panorama management server to use the **Eth1** and **Eth2** interfaces, configure these interfaces on the Log Collector also. For each interface, select the corresponding tab and configure one or both of the following field sets based on the IP protocols of your network. • IPv4—**IP Address**, **Netmask**, and **Default Gateway** • IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway** 5. (Optional) For the **Management** interface, select **SNMP** if you will use an SNMP manager to monitor Log Collector statistics. Using SNMP requires additional steps besides configuring the Log Collector (see Monitor Panorama and Log Collector Statistics Using SNMP). 6. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again. This step is required before you can enable logging disks or assign the Eth1 and Eth2 interfaces to logging functions. 7. (Optional) Select **Panorama > Managed Collectors**, edit the Log Collector, and select the interfaces (MGT, Eth1, or Eth2) that it will use for **Device Log Collection** and **Collector Group Communication** (default is MGT). 8. Select **Disks** and **Add** each RAID disk pair. 9. Click **OK** and **Commit**, set the **Commit Type** to **Panorama**, and click **Commit** again.
Step 5 Add a Firewall as a Managed Device.	Use the web interface of the primary Panorama to perform this task for each firewall that will forward logs to the Log Collectors.

Deploy Panorama with Default Log Collectors (Continued)	
Step 6 Edit the default Collector Group that is predefined on the primary Panorama.	Use the web interface of the primary Panorama to Configure a Collector Group: 1. Select **Panorama > Collector Groups** and select the **default** Collector Group. If you add multiple Log Collectors to a single Collector group, selecting the **Enable log redundancy across collectors** check box is a best practice. 2. (Optional) Select the **Monitoring** tab and configure the settings if you will use SNMP to monitor Log Collectors. 3. Select the **Device Log Forwarding** tab. The Collector Group Members section displays the local Log Collector of the primary Panorama because it is pre-assigned to the default Collector Group. If this Collector Group will contain multiple Log Collectors, assign the local Log Collector of the secondary Panorama. 4. In the Log Forwarding Preferences section, assign firewalls according to the number of Log Collectors in this Collector Group: • Single—Assign the firewalls that will forward logs to the default Log Collector of the primary Panorama, as illustrated in Figure: Single Default Log Collector Per Collector Group. • Multiple—Assign each firewall to both Log Collectors for redundancy. When you configure the preferences, make Log Collector 1 the first priority for half the firewalls and make Log Collector 2 the first priority for the other half, as illustrated in Figure: Multiple Default Log Collectors Per Collector Group. 5. Click **OK** to save your changes.
Step 7 Configure a Collector Group that contains the Log Collector on the secondary Panorama. Required if each Collector Group has only one Log Collector.	Use the web interface of the primary Panorama to Configure a Collector Group: 1. Select **Panorama > Collector Groups** and **Add** the Collector Group. 2. Enter a **Name** to identify the Collector Group. 3. (Optional) Select the **Monitoring** tab and configure the settings if you will use an SNMP manager to monitor Log Collector statistics and traps. 4. Select the **Device Log Forwarding** tab and, in the Collector Group Members section, **Add** the default Log Collector of the secondary Panorama. 5. In the Log Forwarding Preferences section, assign the firewalls that will forward logs to the default Log Collector of the secondary Panorama, as illustrated in Figure: Single Default Log Collector Per Collector Group. 6. Click **OK** to save your changes.

 © Palo Alto Networks, Inc.

Deploy Panorama with Default Log Collectors (Continued)		
Step 8	Commit your changes.	Use the web interface of the primary Panorama to perform the following steps: 1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 2. Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Groups you added, and click **Commit** again.
Step 9	Manually fail over so that the secondary Panorama becomes active.	Use the web interface of the primary Panorama to perform the following steps: 1. Select **Panorama > High Availability**. 2. Click **Suspend local Panorama** in the Operational Commands section.
Step 10	On the secondary Panorama, configure the network settings of the Log Collector that is local to the primary Panorama.	Use the web interface of the secondary Panorama to perform the following steps: 1. In the Panorama web interface, select **Panorama > Managed Collectors** and select the Log Collector that is local to the primary Panorama. 2. Enter the IP address or FQDN of the secondary Panorama in the **Panorama Server IP** field. Enter the IP address or FQDN of the primary Panorama in the **Panorama Server IP 2** field. These fields are required. 3. Select the **Management** tab and complete one or both of the following field sets (depending on the IP protocols of your network) with the management interface values of the primary Panorama: • IPv4—**IP Address**, **Netmask**, and **Default Gateway** • IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway** 4. Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. Wait until the HA synchronization finishes before proceeding. 5. Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Groups you added, and click **Commit** again.
Step 11	Manually fail back so that the primary Panorama becomes active.	Use the web interface of the secondary Panorama to perform the following steps: 1. Select **Panorama > High Availability**. 2. Click **Suspend local Panorama** in the Operational Commands section.

Deploy Panorama with Default Log Collectors (Continued)	
Step 12 Configure log forwarding.	Use the web interface of the primary Panorama to: 1. Configure Log Forwarding to Panorama. 2. Verify Log Forwarding to Panorama. 3. (Optional) Configure Log Forwarding from Panorama to External Destinations. You can configure separate external server profiles for different log types and severity levels. You can also assign separate profiles to each Panorama HA peer. For example, you might want each peer to forward logs to a different syslog server. To make each Panorama peer forward logs to different external services, log in to the web interface of each peer, select **Panorama > Collector Groups**, select the Collector Group, select **Collector Log Forwarding**, assign the server profiles, and click **OK**.

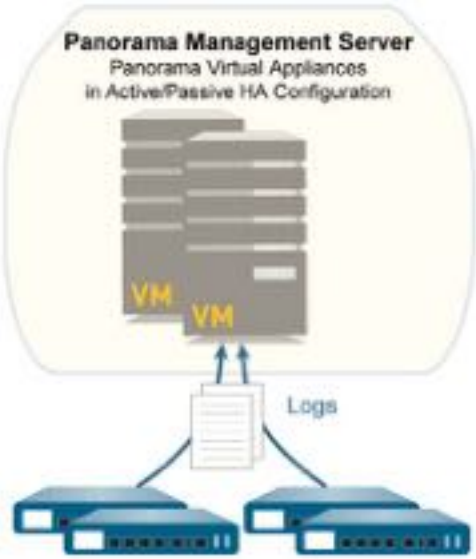
Deploy Panorama Virtual Appliances with Local Log Collection

The following figure illustrates Panorama in a centralized log collection deployment. In this example, the Panorama management server comprises two Panorama virtual appliances that are deployed in an active/passive high availability (HA) configuration. This configuration suits firewall management within a VMware virtual infrastructure in which Panorama processes up to 10,000 logs/second. (For details on deployment options, see Plan a Log Collection Deployment.) The firewalls send logs to the Panorama management server (to its virtual disk or Network File System [NFS] datastore). By default, the active and passive peers both receive logs, though you can Modify Log Forwarding and Buffering Defaults so that only the active peer does. By default, the Panorama virtual appliance uses approximately 11GB on its internal disk partition for log storage, though you can Expand Log Storage Capacity on the Panorama Virtual Appliance if necessary.

If the logging rate increases beyond 10,000 logs per second, it is recommended that you add Dedicated Log Collectors (M-Series appliances in Log Collector mode). Deploy Panorama with Dedicated Log Collectors describes a deployment with Dedicated Log Collectors managed by Panorama virtual appliances or by M-Series appliances in Panorama mode.

 © Palo Alto Networks, Inc.

Figure: Panorama Virtual Appliances with Local Log Collection



Perform the following steps to deploy Panorama virtual appliances with local log collection. Skip any steps you have already performed (for example, the initial setup).

Deploy Panorama Virtual Appliances with Local Log Collection		
Step 1	Perform the initial setup of each Panorama virtual appliance.	1. Install the Panorama Virtual Appliance. 2. Perform Initial Configuration of the Panorama Virtual Appliance. 3. Register Panorama and Install Licenses. 4. Install Content and Software Updates for Panorama. 5. Set Up HA on Panorama.
Step 2	Perform the following steps to prepare Panorama for log collection.	1. Add a Firewall as a Managed Device for each one that will forward logs to Panorama. 2. Configure Log Forwarding to Panorama. 3. (Optional) Configure Log Forwarding from Panorama to External Destinations.
Step 3	Commit your changes.	Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

 © Palo Alto Networks, Inc.

Manage Licenses and Updates

You can use Panorama to centrally manage licenses, software updates, and content updates on firewalls and Dedicated Log Collectors (M-Series appliances in Log Collector mode). When you deploy licenses or updates, Panorama checks in with the Palo Alto Networks licensing server or update server, verifies the request validity, and then allows retrieval and installation of the license or update. This capability facilitates deployment by eliminating the need to repeat the same tasks on each firewall or Dedicated Log Collector. It is particularly useful for managing firewalls that don't have direct Internet access or for managing Dedicated Log Collectors, which don't have a web interface.



Before deploying updates, see Panorama, Log Collector, and Firewall Version Compatibility for important details about update version compatibility among Palo Alto Networks devices.

You must activate a support subscription directly on each firewall; you cannot use Panorama to deploy support subscriptions.

To activate licenses or install updates on the Panorama management server, see Register Panorama and Install Licenses and Install Content and Software Updates for Panorama.

- Manage Licenses on Firewalls Using Panorama
- Deploy Updates to Firewalls and Log Collectors Using Panorama

Manage Licenses on Firewalls Using Panorama

The following steps describe how to retrieve new licenses using an auth-code and push the license keys to managed firewalls. It also describes how to manually update (refresh) the license status of firewalls that do not have direct Internet access. For firewalls that have direct Internet access, Panorama automatically performs a daily check-in with the licensing server, retrieves license updates and renewals, and pushes them to the firewalls. The check-in is hard-coded to occur between 1 and 2 A.M.; you cannot change this schedule.



You cannot use Panorama to activate the support license of firewalls. You must access the firewalls individually to activate their support licenses.
To activate licenses for Panorama itself, see Register Panorama and Install Licenses.

Manage Licenses on Firewalls Using Panorama	
• Activate newly purchased licenses.	1. Select **Panorama > Device Deployment > Licenses** and click **Activate**. 2. Enter the **Auth Code** that Palo Alto Networks provided for each firewall that has a new license. 3. Click **Activate**. 4. (WildFire subscriptions only) Perform a commit on each firewall that has a new WildFire subscription to complete the activation: • Commit any pending changes. You must access each firewall web interface to do this. • Make a minor change and perform a commit. For example, update a rule description and commit the change. If the firewalls belong to the same device group, you can push the rule change from Panorama to initiate a commit on all those firewalls instead of accessing each firewall separately.  Check that the WildFire Analysis profile rules include the advanced file types that the WildFire subscription supports. If no rule changes are required, make a minor edit to a rule description and perform a commit.
• Update the license status of firewalls.	1. Select **Panorama > Device Deployment > Licenses**. Each entry on the page indicates whether the license is active or inactive and displays the expiration date for active licenses. 2. If you previously activated auth-codes for the support subscription directly on the firewalls, click **Refresh** and select the firewalls from the list. Panorama retrieves the license, deploys it to the firewalls and updates the licensing status on the Panorama web interface.

 © Palo Alto Networks, Inc.

Deploy Updates to Firewalls and Log Collectors Using Panorama

You can use Panorama to qualify software and content updates by deploying them to a subset of firewalls or Dedicated Log Collectors before installing the updates on the rest. If you want to schedule periodic content updates, Panorama requires a direct Internet connection. To deploy software or content updates on demand (unscheduled), the procedure differs based on whether Panorama is connected to the Internet.



By default, you can download up to two software or content updates of each type to Panorama. When you start any download beyond that maximum, Panorama deletes the oldest update of the selected type. To change the maximum, see Manage Panorama Storage for Software and Content Updates.

Panorama displays a warning if you manually deploy a content update when a scheduled update process has started or will start within five minutes.

- Supported Updates
- Schedule a Content Update Using Panorama
- Deploy an Update to Log Collectors when Panorama is Internet-connected
- Deploy an Update to Log Collectors when Panorama is not Internet-connected
- Deploy an Update to Firewalls when Panorama is Internet-connected
- Deploy an Update to Firewalls when Panorama is not Internet-connected

Supported Updates

The software and content updates you can install vary based on which subscriptions are active on each firewall or Log Collector:

Platform Type	Software Updates	Content Updates
Log Collector	Panorama	Applications (Log Collectors don't need Threats signatures) Antivirus BrightCloud URL filtering WildFire
Firewall	PAN-OS GlobalProtect agent/app	Applications Applications and Threats Antivirus BrightCloud URL filtering WildFire GlobalProtect data files

Schedule a Content Update Using Panorama

Panorama requires a direct Internet connection for scheduling Supported Updates on firewalls and Log Collectors. Otherwise, you can perform only on-demand updates. (To schedule Antivirus, WildFire, or BrightCloud URL updates for Log Collectors, they must run Panorama 7.0.3 or a later release.) Each firewall or Log Collector receiving an update generates a log to indicate that the installation succeeded (Config log) or failed (System log). You cannot schedule content updates for the Panorama management server; to install updates on demand, see Install Content and Software Updates for Panorama.

Before deploying updates, see Panorama, Log Collector, and Firewall Version Compatibility for important details about update version compatibility. Refer to the Release Notes for the minimum content release version you must install for a Panorama release.

Panorama can download only one update at a time; stagger the updates to ensure they succeed. If you schedule multiple updates to download during the same time interval, only the first download will succeed.

Perform the following steps for each update type you want to schedule.

Schedule a Content Update Using Panorama	
Step 1	Select **Panorama > Device Deployment > Dynamic Updates**, click **Schedules**, and click **Add**.
Step 2	Specify a **Name** to identify the schedule, the update **Type**, and the update frequency (**Recurrence**). The available frequency options depend on the update **Type**. PAN-OS uses the Panorama timezone for update scheduling. If you set the **Type** to **App and Threat**, Log Collectors install and need only the Applications content, not the Threats content. Firewalls use both Applications and Threats content. For details, see Panorama, Log Collector, and Firewall Version Compatibility. The WildFire Private (**WF-Private**) type is available only if you set the **WildFire Server** field (**Panorama > Setup > WildFire**) to a WF-500 appliance, not to the WildFire cloud.
Step 3	Select one of the following schedule actions and then select the firewalls or Log Collectors: • **Download And Install** (best practice)—Select **Devices** (firewalls) or **Log Collectors**. • **Download Only**—Panorama downloads the update but does not install it.
Step 4	Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

Deploy an Update to Log Collectors when Panorama is Internet-connected

For a list of software and content updates you can install on Log Collectors, see Supported Updates.

Deploy an Update to Log Collectors when Panorama is Internet-connected		
Step 1	Upgrade the Panorama management server to software and content versions that are the same as or higher than the versions you will install on the Log Collectors.	Install Content and Software Updates for Panorama. You must upgrade Panorama and then Log Collectors before upgrading firewalls.

 © Palo Alto Networks, Inc.

Deploy an Update to Log Collectors when Panorama is Internet-connected (Continued)		
Step 2	Install content updates.  You must install content updates before software updates. Refer to the Release Notes for the minimum content version you must install for a Panorama release.	1. Select **Panorama > Device Deployment > Dynamic Updates**. 2. Click **Check Now** to display the latest updates. If an update is available, the Action column displays **Download**. 3. Review the Version and File Name columns to determine which update you want to install. 4. Install the Applications or Applications and Threats update first, and then install any other updates (Antivirus, WildFire, or URL Filtering) one at a time in any sequence.  Regardless of whether your subscription includes both Applications and Threats content, Panorama installs and needs only the Applications content. For details, see Panorama, Log Collector, and Firewall Version Compatibility. a. Click **Download** in the Action column for the desired update. After a successful download, the Action value changes to **Install**. b. Click **Install**, select the Log Collectors, and click **OK**.
Step 3	Determine the software upgrade path for the Log Collectors. Required for Panorama software updates.	Select **Panorama > Managed Collectors** and note the current Software Version for the Log Collectors you will upgrade. You cannot skip any major Panorama release versions on the path between the current version and the target version. For example, to upgrade Log Collectors from Panorama 5.0.13 to Panorama 7.0.3: 1. Download and install a Panorama 5.1.x release based on your platform: • Panorama virtual appliance—Download and install Panorama 5.1.0 and reboot. • Panorama M-Series appliance: – Download Panorama 5.1.0 without installing or rebooting. – Download and install a later Panorama 5.1.x release and reboot. 2. Download and install Panorama 6.0.0 and reboot. 3. Download and install Panorama 6.1.0 and reboot. 4. Download Panorama 7.0.1 and upload it to the Log Collectors without installing or rebooting (7.0.1 is the base image for the 7.0 release; not 7.0.0). 5. Download and install Panorama 7.0.3 and reboot.

Deploy an Update to Log Collectors when Panorama is Internet-connected (Continued)		
Step 4	Install software updates.	1. Select **Panorama > Device Deployment > Software**. 2. Click **Check Now** to display the latest updates. If an update is available, the Action column displays **Download**. 3. Review the Version, File Name, and Platform columns to determine which update to deploy. For Panorama software, the Platform column displays `m`. 4. Click **Download** in the Action column for the desired update. After a successful download, the Action value changes to **Install**. 5. Click **Install** and select the Log Collectors. 6. Select one of the following based on the update version you are installing within the upgrade path (Step 3): • **Upload only to device (do not install)** • **Reboot device after install** 7. Click **OK** to start the installation or upload.

Step 5	Verify the software and content update versions that are installed on the Log Collector.	Log in to the Log Collector CLI and enter the **`show system info`** operational command. The output will resemble the following:
		`sw-version: 7.0.3`
		`app-version: 497-1738`
		`app-release-date: 2015/10/29 15:46:03`
		`av-version: 1168-1550`
		`av-release-date: 2015/10/21 14:31:27`
		`threat-version: 497-1738`
		`threat-release-date: 2015/10/29 15:46:03`

Deploy an Update to Log Collectors when Panorama is not Internet-connected

For a list of software and content updates you can install on Log Collectors, see Supported Updates.

Deploy an Update to Log Collectors when Panorama is not Internet-connected		
Step 1	Upgrade the Panorama management server to software and content versions that are the same as or higher than the versions you will install on the Log Collectors.	Install Content and Software Updates for Panorama. You must upgrade Panorama and then Log Collectors before upgrading firewalls.

 © Palo Alto Networks, Inc.

Deploy an Update to Log Collectors when Panorama is not Internet-connected (Continued)		
Step 2	Determine the software upgrade path. Required for Panorama software updates.	Log in to Panorama, select **Panorama > Managed Collectors**, and note the current Software Version for the Log Collectors you will upgrade. You cannot skip any major Panorama release versions on the path between the current version and the target version. For example, to upgrade Log Collectors from Panorama 5.0.13 to Panorama 7.0.3: 1. Upload and install a Panorama 5.1.x release based on your platform: • Panorama virtual appliance—Upload and install Panorama 5.1.0 and reboot. • Panorama M-Series appliance: – Upload Panorama 5.1.0 without installing or rebooting. – Upload and install a later Panorama 5.1.x release and reboot. 2. Upload and install Panorama 6.0.0 and reboot. 3. Upload and install Panorama 6.1.0 and reboot. 4. Upload Panorama 7.0.1 to the Log Collector without installing or rebooting (7.0.1 is the base image for the 7.0 release; not 7.0.0). 5. Upload and install Panorama 7.0.3 and reboot.
Step 3	Download the updates to a host that has Internet access. Panorama must have access to the host.	1. Use a host with Internet access to log in to the Palo Alto Networks Customer Support web site. 2. Download content updates: a. Click **Dynamic Updates** in the Resources section. b. **Download** the desired content update and save the file to the host. Perform this step for each content type you will update. 3. Download software updates: a. Return to the main page of the Palo Alto Networks Customer Support web site and click **Software Updates** in the Resources section. b. Review the Download column to determine the version to install. The filename of the update package indicates the platform: Panorama-ESX-<release> for the Panorama virtual appliance (VMware ESXi server or vCloud Air) or Panorama-m-<release> for the Panorama M-Series appliance. c. Click the filename and save the file to the host.

Deploy an Update to Log Collectors when Panorama is not Internet-connected (Continued)	
Step 4 Install content updates. You must install content updates before software updates. Refer to the Release Notes for the minimum content release version you must install for a Panorama release.	Install the Applications or Applications and Threats update first, and then install any other updates (Antivirus, WildFire, or URL Filtering) one at a time in any sequence. Regardless of whether your subscription includes both Applications and Threats content, Panorama installs and needs only the Applications content. For details, see Panorama, Log Collector, and Firewall Version Compatibility. 1. In Panorama, select **Panorama > Device Deployment > Dynamic Updates**. 2. Click **Upload**, select the update **Type**, **Browse** to the update file on the host, and click **OK**. 3. Click **Install From File**. 4. Select the update **Type** and select the **File Name** of the update you just uploaded. 5. Select the Log Collectors. 6. Click **OK** to start the installation.
Step 5 Install software updates.	1. Select **Panorama > Device Deployment > Software**. 2. Click **Upload**, **Browse** to the update file on the host, and click **OK**. 3. Click **Install** in the Action column. 4. Select the Log Collectors on which to install the update. 5. Select one of the following based on the update version you are installing within the upgrade path (Step 2): • **Upload only to device (do not install)** • **Reboot device after install** 6. Click **OK** to start the installation.

Step 6 Verify the software and/or content versions that are installed on each Log Collector.	Log in to the Log Collector CLI and enter the **`show system info`** operational command. The output will resemble the following:
	`sw-version: 7.0.3`
	`app-version: 497-1738`
	`app-release-date: 2015/10/29 15:46:03`
	`av-version: 1168-1550`
	`av-release-date: 2015/10/21 14:31:27`
	`threat-version: 497-1738`
	`threat-release-date: 2015/10/29 15:46:03`

Deploy an Update to Firewalls when Panorama is Internet-connected

Before deploying updates to firewalls, you must upgrade Panorama and then upgrade Log Collectors. For a list of software and content updates you can install on firewalls, see Supported Updates.

You cannot deploy GlobalProtect data file updates on demand; you can only schedule the updates using the firewall web interface or a Panorama template.

 © Palo Alto Networks, Inc.

Deploy an Update to Firewalls when Panorama is not Internet-connected

For a list of software and content updates you can install on firewalls, see Supported Updates.

Deploy an Update to Firewalls when Panorama is not Internet-connected		
Step 1	Upgrade the Panorama management server and Log Collectors to software and content versions that are the same as or higher than the versions you will install on the firewalls. For important software and content compatibility details, see Panorama, Log Collector, and Firewall Version Compatibility.	You must upgrade Panorama and then Log Collectors before upgrading firewalls. 1. Install Content and Software Updates for Panorama. 2. Deploy an Update to Log Collectors when Panorama is not Internet-connected.
Step 2	Save a backup of the current configuration file of each firewall for which you will upgrade software. You can use this backup to restore the configuration if you have problems with the upgrade. Although the firewall automatically creates a configuration backup, the best practice is to create and externally store a backup before upgrading.	1. Log in to Panorama and select **Panorama > Setup > Operations**. 2. Click **Save named Panorama configuration snapshot**, enter a **Name** for the configuration, and click **OK**. 3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, click **OK**, and save the exported file.
Step 3	Determine the software upgrade path. Required for PAN-OS software updates.	Select **Panorama > Managed Devices**, and note the current Software Version for the firewalls you will upgrade. You cannot skip any major PAN-OS release versions on the path between the current version and the target version. For example, to upgrade firewalls from PAN-OS 5.0.13 to PAN-OS 7.0.3: 1. Upload and install PAN-OS 5.1.0 and reboot. 2. Upload and install PAN-OS 6.0.0 and reboot. 3. Upload and install PAN-OS 6.1.0 and reboot. 4. Upload PAN-OS 7.0.1 to the firewalls without installing or rebooting (7.0.1 is the base image for the 7.0 release; not 7.0.0). 5. Upload and install PAN-OS 7.0.3 and reboot.

Deploy an Update to Firewalls when Panorama is not Internet-connected (Continued)		
Step 4	Download the updates to a host that has Internet access. Panorama must have access to the host.	1. Use a host with Internet access to log in to the Palo Alto Networks Customer Support web site. 2. Download content updates: a. Click **Dynamic Updates** in the Resources section. b. **Download** the desired content update and save the file to the host. Perform this step for each content type you will update. 3. Download software updates: a. Return to the main page of the Palo Alto Networks Customer Support web site and click **Software Updates** in the Resources section. b. Review the Download column to determine the version to install. The filename of the update package indicates the platform. c. Click the filename and save the file to the host.
Step 5	Install content updates. You must install content updates before software updates. To see the minimum content release version you must install for a PAN-OS release, refer to the Release Notes.	Install the Applications or Applications and Threats update first, and then install any other updates (Antivirus, WildFire, or URL Filtering) one at a time in any sequence. 1. Select **Panorama > Device Deployment > Dynamic Updates**. 2. Click **Upload**, select the update **Type**, **Browse** to the update file, and click **OK**. 3. Click **Install From File**. 4. Select the update **Type** and select the **File Name** of the content update you just uploaded. 5. Select the firewalls on which to install the update. 6. Click **OK** to start the installation.
Step 6	Activate a GlobalProtect agent/app software update on firewalls. You activate the update on firewalls so that users can download it to their client systems.	1. Select **Panorama > Device Deployment > GlobalProtect Client**. 2. Click **Upload**, **Browse** to the update file, and click **OK**. 3. Click **Activate From File**. 4. Select the **File Name** of the GlobalProtect agent/app update you just uploaded. 5. Select the firewalls on which to install the update. 6. Click **OK** to start the installation.
Step 7	Upload PAN-OS software updates.	1. Select **Panorama > Device Deployment > Software**. 2. Click **Upload**, **Browse** to the update file, and click **OK**.

 © Palo Alto Networks, Inc.

Deploy an Update to Firewalls when Panorama is not Internet-connected (Continued)		
Step 8	Install PAN-OS software updates.  To avoid downtime when updating the software on high availability (HA) firewalls, update one HA peer at a time. For active/active firewalls, it doesn't matter which peer you update first. For active/passive firewalls, you must update the passive peer first, suspend the active peer (fail over), update the active peer, and then return the active peer to a functional state (fail back).	Perform the steps that apply to your firewall deployment. Remember that rebooting is necessary only for certain update versions within the upgrade path (see Step 3). • **Non-HA firewalls**—Click **Install** in the Action column, select all the firewalls you are upgrading, select **Reboot device after install**, and click **OK**. • **Active/active HA firewalls:** a. Click **Install**, clear **Group HA Peers**, select either HA peer, select **Reboot device after install**, and click **OK**. Wait for the firewall to finish rebooting before proceeding. b. Click **Install**, clear **Group HA Peers**, select the HA peer that you didn't update yet, select **Reboot device after install**, and click **OK**. • **Active/passive HA firewalls**—In this example, the active firewall is named fw1 and the passive firewall is named fw2: a. Click **Install**, clear **Group HA Peers**, select fw2, select **Reboot device after install**, and click **OK**. Wait for fw2 to finish rebooting before proceeding. b. Access fw1, select **Device > High Availability > Operational Commands**, and click **Suspend local device**. c. Access fw2 and, on the **Dashboard**, **High Availability** widget, verify that the **Local** firewall state is **active** and the **Peer** firewall is **suspended**. d. Access Panorama, select **Panorama > Device Deployment > Software**, click **Install**, clear **Group HA Peers**, select fw1, select **Reboot device after install**, and click **OK**. Wait for fw1 to finish rebooting before proceeding. e. Access fw1, select **Device > High Availability > Operational Commands**, and click **Make local device functional**. Wait two minutes before proceeding. f. On fw1, select the **Dashboard** tab and, in the **High Availability** widget, verify that the **Local** firewall state is **active** and the **Peer** firewall is **passive**.
Step 9	Verify the software and content versions that are installed on each managed firewall.	1. Select **Panorama > Managed Devices**. 2. Locate the firewall and review the values in the Software Version, Apps and Threat, Antivirus, URL Filtering, and GlobalProtect Client columns.

 © Palo Alto Networks, Inc.

Monitor Network Activity

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

This section covers the following topics:

- Use Panorama for Visibility
- Use Case: Monitor Applications Using Panorama
- Use Case: Respond to an Incident Using Panorama

Use Panorama for Visibility

In addition to its central deployment and firewall configuration features, Panorama also allows you to monitor and report on all traffic that traverses your network. While the reporting capabilities on Panorama and the firewall are very similar, the advantage that Panorama provides is that it is a single pane view of aggregated information across all your managed firewalls. This aggregated view provides actionable information on trends in user activity, traffic patterns, and potential threats across your entire network.

Using the Application Command Center (ACC), the App-Scope, the log viewer, and the standard, customizable reporting options on Panorama, you can quickly learn more about the traffic traversing the network. The ability to view this information allows you to evaluate where your current policies are adequate and where they are insufficient. You can then use this data to augment your network security strategy. For example, you can enhance the security rules to increase compliance and accountability for all users across the network, or manage network capacity and minimize risks to assets while meeting the rich application needs for the users in your network.

The following topics provide a high-level view of the reporting capabilities on Panorama, including a couple of use cases to illustrate how you can use these capabilities within your own network infrastructure. For a complete list of the available reports and charts and the description of each, refer to the online help.

▲ Monitor the Network with the ACC and AppScope

▲ Analyze Log Data

▲ Generate, Schedule, and Email Reports

Monitor the Network with the ACC and AppScope

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

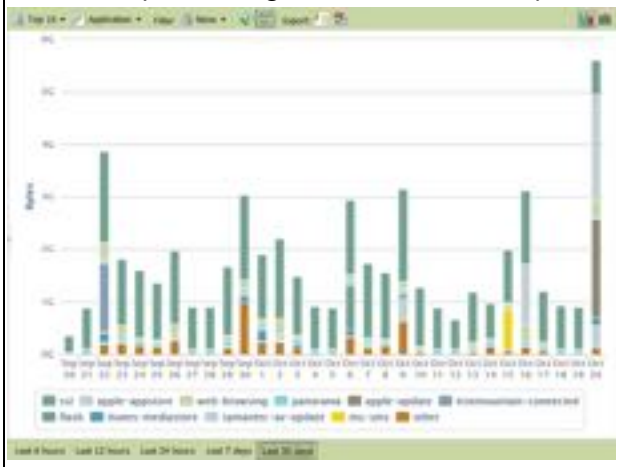
The ACC on Panorama displays a summary of network traffic. Panorama can dynamically query data from all the managed firewalls on the network and display it in the ACC. This display allows you to monitor the traffic by applications, users, and content activity—URL categories, threats, security policies that effectively block data or files—across the entire network of Palo Alto Networks next-generation firewalls.

The AppScope helps identify unexpected or unusual behavior on the network at a glance. It includes an array of charts and reports—Summary Report, Change Monitor, Threat Monitor, Threat Map, Network Monitor, Traffic Map—that allow you to analyze traffic flows by threat or application, or by the source or destination for the flows. You can also sort by session or byte count.

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

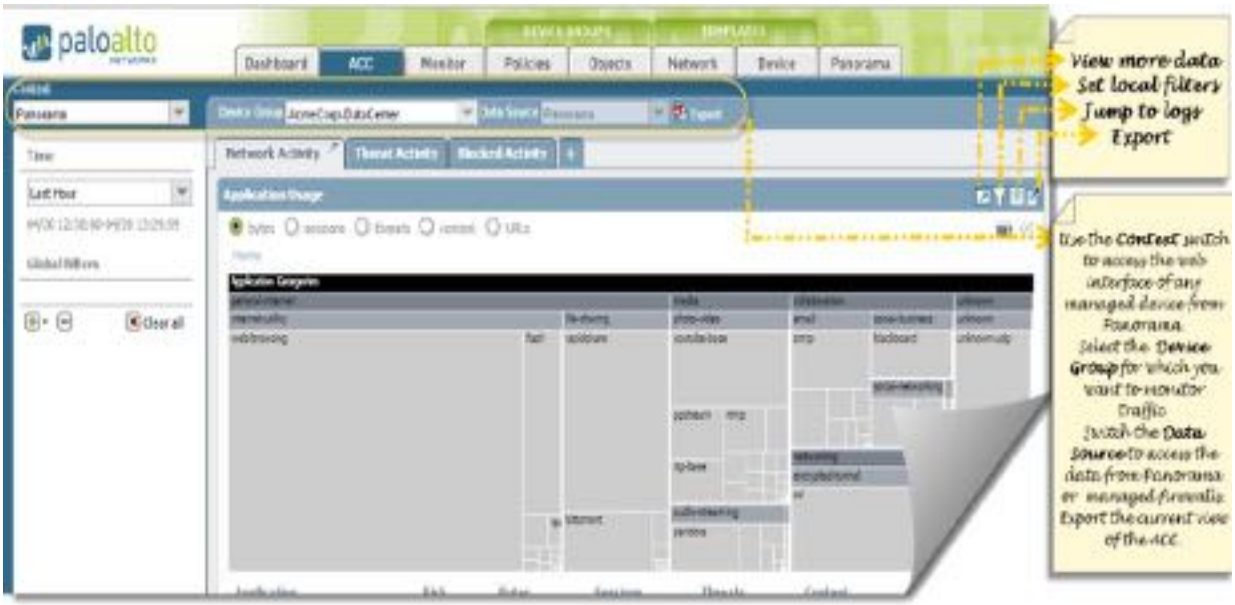
ACC	Monitor > AppScope
• What are the top applications used on the network and how many are high-risk applications? Who are the top users of high-risk applications on the network? • What are the top URL categories being viewed in the last hour?	• What are the Application usage trends—what are the top five applications that have gained use and the top five that have decreased in use? • How has user activity changed over the current week as compared to last week or last month?

 © Palo Alto Networks, Inc.

ACC	Monitor > AppScope
• What are the top bandwidth-using applications? Who are the users/hosts that consume the highest bandwidth? • What content or files are being blocked and are there specific users who trigger this File Blocking/Data Filtering rule? • What is the amount of traffic exchanged between two specific IP addresses or generated by a specific user? Where is the destination server or client located geographically?	• Which users and applications take up most of the network bandwidth? And how has this consumption changed over the last 30 days?  • What are the threats on the network, and how are these incoming and outgoing traffic threats distributed geographically?

You can then use the information to maintain or enforce changes to the traffic patterns on your network. See Use Case: Monitor Applications Using Panorama for a glimpse into how the visibility tools on Panorama can influence how you shape the acceptable use policies for your network.

Here are a few tips to help you navigate the ACC:



- **Switch from a Panorama view to a Device view**—Use the **Context** drop-down to access the web interface of any managed firewall. For details, see Context Switch—Firewall or Panorama.

- **Change Device Group and Data Source**—The default **Data Source** used to display the statistics on the charts in the ACC is **Panorama** local data, and the default **Device Group** setting is **All**. Using the local data on Panorama provides a quick load time for the charts. You can, however, change the data source to **Remote Device Data** if all the managed devices are on PAN-OS 7.0. If the managed firewalls have a mix of PAN-OS 7.0 and earlier releases, you can only view Panorama data. When configured to use Remote Device Data, Panorama will poll all the managed firewalls and present an aggregated view of the data. The onscreen display indicates the total number of firewalls being polled and the number of firewalls that have responded to the query for information.
- **Select the Tabs and Widgets to View**—The ACC includes three tabs and an array of widgets that allow you to find the information that you care about. With the exception of the application usage widget and host information widget, all the other widgets display data only if the corresponding feature has been licensed on the firewall, and you have enabled logging.
- **Tweak Time Frame and Refine Data**—The reporting time period in the ACC ranges from the last 15 minutes to the last hour, day, week, month, or any custom-defined time. By default, each widget displays the top 10 items and aggregates all the remaining items as **others**. You can sort the data in each widget using various attributes—for example, sessions, bytes, threats, content, and URLs. You can also set local filters to filter the display within the table and graph in a widget, and then promote the widget filter as a global filter to pivot the view across all the widgets in the ACC.

Analyze Log Data

The **Monitor** tab on Panorama provides access to log data; these logs are an archived list of sessions that have been processed by the managed firewalls and forwarded to Panorama.

Log data can be broadly grouped into two types: those that detail information on traffic flows on your network such as applications, threats, host information profiles, URL categories, content/file types and those that record system events, configuration changes and alarms.

Based on the log forwarding configuration on the managed firewalls, the **Monitor > Logs** tab can include logs for traffic flows, threats, URL filtering, data filtering, host information profile (HIP) matches, and WildFire submissions. You can review the logs to verify a wealth of information on a given session or transaction. Some examples of this information are the user who initiated the session, the action (allow or deny) that the firewall performed on the session, and the source and destination ports, zones, and addresses. The System and Config logs can indicate a configuration change or an alarm that the firewall triggered when a configured threshold was exceeded.



If Panorama will manage firewalls running software versions earlier than PAN-OS 7.0, specify a WildFire server from which Panorama can gather analysis information for WildFire samples that those firewalls submit. Panorama uses the information to complete WildFire Submissions logs that are missing field values introduced in PAN-OS 7.0. Firewalls running earlier releases won't populate those fields. To specify the server, select **Panorama > Setup > WildFire**, edit the General Settings, and enter the **WildFire Server** name. The default is **wildfire-public-cloud**, which is the WildFire cloud hosted in the United States.

 © Palo Alto Networks, Inc.

Generate, Schedule, and Email Reports

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



It is recommended that you install matching software releases on Panorama and the firewalls for which you will generate reports. For example, if the Panorama management server runs Panorama 6.1, install PAN-OS 6.1 on its managed firewalls before generating the reports. This practice avoids issues that might occur if you create reports that include fields supported in the Panorama release but not supported in an earlier PAN-OS release on the firewalls.

You can create the following types of reports:

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

Generate, Schedule, and Email Reports	
Step 1 Generate reports.	The steps to generate a report depend on the type: • Create a custom report. a. Select **Monitor > Manage Custom Reports**. b. Click **Add** and enter a **Name** for the report. c. Select a **Database** for the report. You can use a summary database or detailed logs on Panorama or on the managed firewalls. d. Select the **Scheduled** check box. e. Define your filtering criteria. Select the **Time Frame**, the **Sort By** order, **Group By** preference, and select the columns that must display in the report. f. (Optional) Select the **Query Builder** attributes to further refine the selection criteria. g. To test the report settings, select **Run Now**. If necessary, modify the settings to change the information that the report displays. h. Click **OK** to save the custom report. • Run a **PDF Summary Report**. a. Select **Monitor > PDF Reports > Manage PDF Summary**. b. Click **Add** and enter a **Name** for the report. c. Use the drop-down for each report group and select one or more of the elements to design the PDF Summary Report. You can include up to 18 elements. d. Click **OK** to save the settings.
Step 2 Configure a **Report Group**. It can include predefined reports, PDF Summary reports, and custom reports. Panorama compiles all the included reports into a single PDF.	1. Select **Monitor > PDF Reports > Report Groups**. 2. Click **Add** and enter a **Name** for the report group. 3. (Optional) Select the **Title Page** check box and add a **Title** for the PDF output. 4. Select from the Predefined Report, PDF Summary Report and the Custom Report lists. 5. Click **Add** to include the selected reports in the report group. 6. Click **OK** to save the settings.

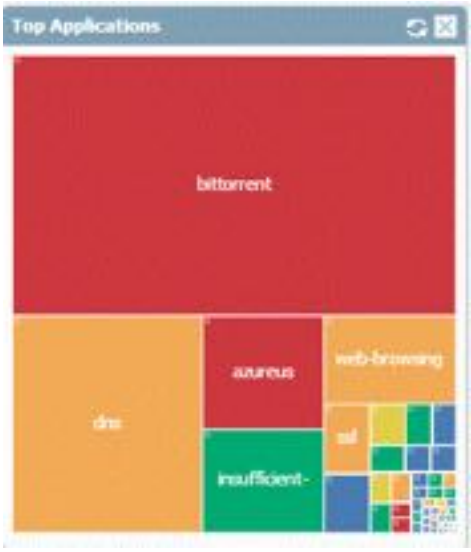
 © Palo Alto Networks, Inc.

Generate, Schedule, and Email Reports (Continued)	
Step 3 Configure an Email server profile.	1. Select **Panorama > Server Profiles > Email**. 2. Click **Add** and enter a **Name** for the profile. 3. For each Simple Mail Transport Protocol (SMTP) server (up to four), click **Add** and enter the information required to connect to the server and send email: • **Name**—A name to identify the SMTP server (1-31 characters). This field is just a label and doesn't have to be the hostname of an existing server. • **Email Display Name**—The name to display in the From field of the email. • **From**—The email address where notification emails will be sent from. • **To**—The email address to which notification emails will be sent. • **Additional Recipient**—To send notifications to a second account, enter the additional address here. • **Email Gateway**—The IP address or hostname of the SMTP gateway to use to send the emails. 4. Click **OK** to save the profile.
Step 4 Schedule the report for email delivery.	1. Select **Monitor > PDF Reports > Email Scheduler**. 2. Click **Add** and enter a **Name** for the email scheduler profile. 3. Select the **Report Group**, the Email server profile you just created (**Email Profile**), and the **Recurrence** for the report. 4. To verify that the email settings are accurate, click **Send test email**. 5. Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

Use Case: Monitor Applications Using Panorama

This example takes you through the process of assessing the efficiency of your current policies and determining where you need to adjust them to fortify the acceptable use policies for your network.

When you log in to Panorama, the **Top Applications** widget on the **Dashboard** gives a preview of the most used applications over the last hour. You can either glance over the list of top applications and mouse over each application block that you want to review the details for, or you can navigate to the **ACC** tab to view the same information as an ordered list. The following image is a view of the **Top Applications** widget on the **Dashboard**.



The data source for this display is the application statistics database; it does not use the traffic logs and is generated whether or not you have enabled logging for security rules. This view into the traffic on your network depicts everything that is allowed on your network and is flowing through unblocked by any policy rules that you have defined.

You can select and toggle the **Data Source** to be local on **Panorama** or you can query the managed firewalls (**Remote Device Data**) for the data; Panorama automatically aggregates and displays the information. For a speedier flow, consider using Panorama as the data source (with log forwarding to Panorama enabled) because the time to load data from the remote devices varies by the time period for which you choose to view data and the volume of traffic that is generated on your network. If your managed firewalls have a combination of PAN-OS 7.0 and earlier versions, remote device data is not available.

Going back to the list of top applications, we can see that bittorrent is very popular. If you now click into the link for the bittorrent application, bittorrent is applied as a global filter and the display in the **ACC** pivots to show information on the application, users who accessed the application, and the details on the risk level and characteristics of the application.

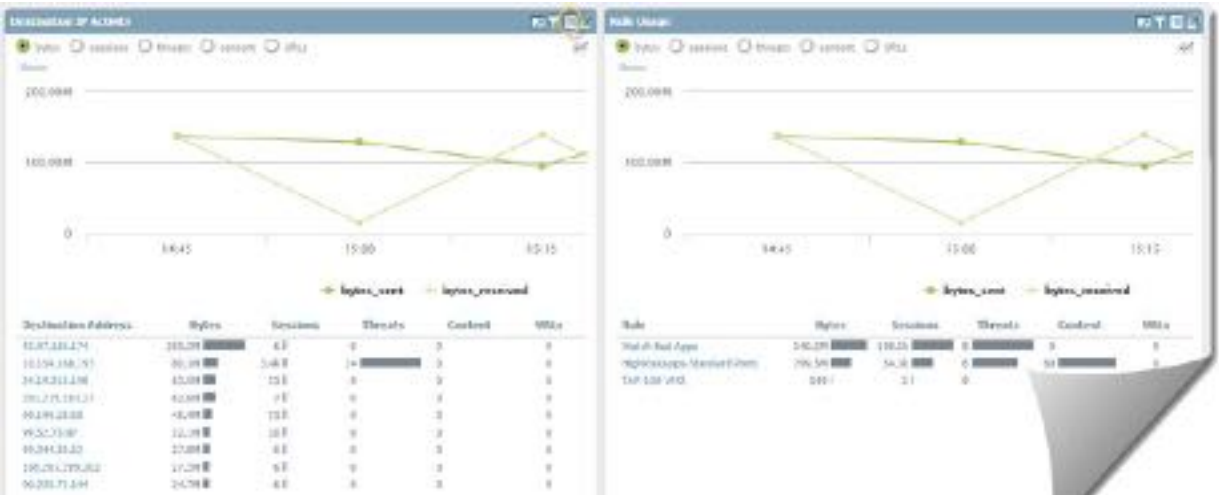
 © Palo Alto Networks, Inc.



In the **User Activity** widget, you can also see how many users are using BitTorrent and the volume of traffic being generated. If you have enabled User-ID, you will be able to view the names of the users who are generating this traffic, and drill in to review all the sessions, content or threats associated with each user. View the **Compromised Hosts** widget in the **Threat Activity** tab to see what correlation objects were matched on, and view the match evidence associated with the user and application. You can also view the threat name, category and ID in the **Threat Activity** widget.



With BitTorrent set as a global filter, use the **Destination IP Activity** and the **Destination Regions** widgets to verify where the traffic was destined. You can also view the ingress and egress zones and the security rule that is letting this connection through.



For more detailed information, jump into the traffic logs for a filtered view and review each log entry for ports used, packets sent, bytes sent and received. Adjust the columns to view more information or less information based on your needs.

The **Monitor > App-Scope> Traffic Map** tab displays a geographical map of the traffic flow and provides a view of incoming versus outgoing traffic. You can also use the **Monitor > App-Scope > Change Monitor** tab to view changes in traffic patterns. For example, compare the top applications used over this hour to the last week or month to determine if there is a pattern or trend.

With all the information you have now uncovered, you can evaluate what changes to make to your policy configurations. Here are some suggestions to consider:

- Be restrictive and to create a pre-rule on Panorama to block all BitTorrent traffic. Then use Panorama device groups to create and push this policy rule to one or more firewalls.
- Enforce bandwidth use limits and create a QoS profile and policy rule that de-prioritizes non-business traffic. Use Panorama device groups and templates to configure QoS and then push rules to one or more firewalls.
- Reduce risk to your network assets and create an application filter that blocks all file sharing applications that are peer-to-peer technology with a risk factor of 4 or 5. Make sure to verify that the bittorrent application is included in that application filter, and will therefore be blocked.
- Schedule a custom report group that pulls together the activity for the specific user and that of top applications used on your network to observe that pattern for another week or two before taking action.

Besides checking for a specific application, you can also check for any unknown applications in the list of top applications. These are applications that did not match a defined App-ID signature and display as *unknown-udp* and *unknown-tcp*. To delve into these unknown applications, click on the name to drill down to the details for the unclassified traffic.

Use the same process to investigate the top source IP addresses of the hosts that initiated the unknown traffic along with the IP address of the destination host to which the session was established. For unknown traffic, the traffic logs, by default, perform a packet capture (pcap) when an unknown application is detected. The green arrow in the left column represents the packet capture snippet of the application data. Clicking on the green arrow displays the pcap in the browser.

Having the IP addresses of the servers (destination IP), the destination port, and the packet captures, you will be better positioned to identify the application and make a decision on how you would like to take action on your network. For example, you can create a custom application that identifies this traffic instead of

 © Palo Alto Networks, Inc.

labeling it as unknown TCP or UDP traffic. Refer to the article Identifying Unknown Applications for more information on identifying unknown application and Custom Application Signatures for information on developing custom signatures to discern the application.

Use Case: Respond to an Incident Using Panorama

Network threats can originate from different vectors, including malware and spyware infections due to drive-by downloads, phishing attacks, unpatched servers, and random or targeted denial of service (DoS) attacks, to name a few methods of attack. The ability to react to a network attack or infection requires processes and systems that alert the administrator to an attack and provide the necessary forensics evidence to track the source and methods used to launch the attack.

The advantage that Panorama provides is a centralized and consolidated view of the patterns and logs collected from the managed firewalls across your network. You can use the information from the automated correlation engine alone or in conjunction with the reports and logs generated from a Security Information Event Manager (SIEM), to investigate how an attack was triggered and how to prevent future attacks and loss of damage to your network.

The questions that this use case probes are:

- How are you notified of an incident?
- How do you corroborate that the incident is not a false positive?
- What is your immediate course of action?
- How do you use the available information to reconstruct the sequence of events that preceded or followed the triggering event?
- What are the changes you need to consider for securing your network?

This use case traces a specific incident and shows how the visibility tools on Panorama can help you respond to the report.

▲ Incident Notification

▲ Review the Widgets in the ACC

▲ Review Threat Logs

▲ Review WildFire Logs

▲ Review Data Filtering Logs

▲ Update Security Rules

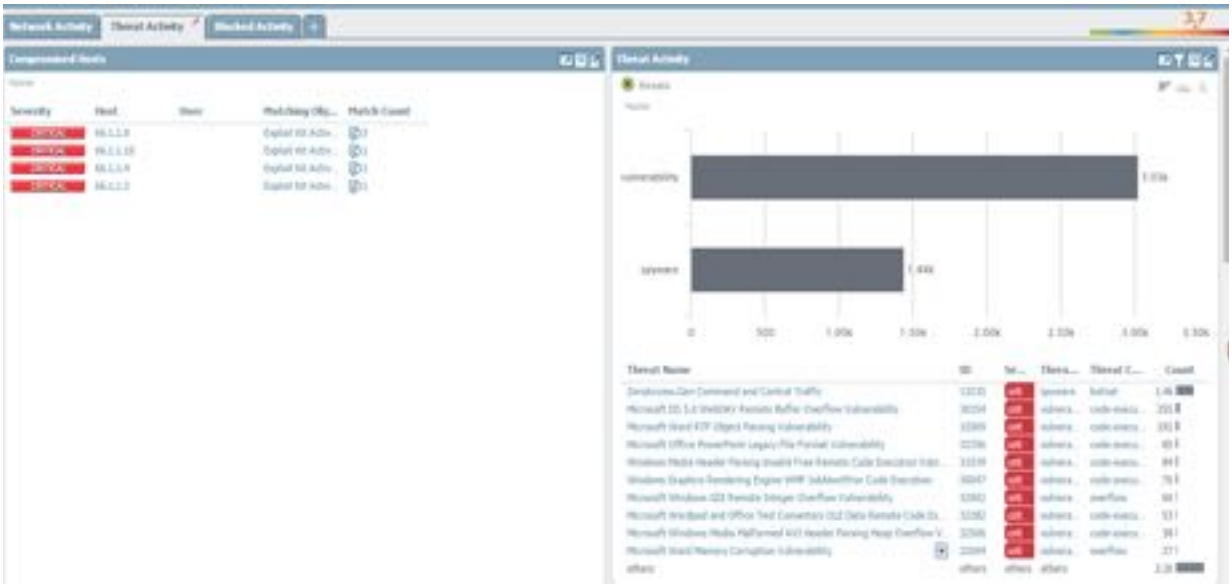
Incident Notification

There are several ways that you could be alerted to an incident depending on how you've configured the Palo Alto Networks firewalls and which third-party tools are available for further analysis. You might receive an email notification that was triggered by a log entry recorded to Panorama or to your syslog server, or you might be informed through a specialized report generated on your SIEM solution, or a third-party paid service or agency might notify you. For this example, let's say that you receive an email notification from Panorama. The email informs you of an event that was triggered by an alert for a **Zero Access gent.Gen Command And Control Traffic** that matched against a spyware signature. Also listed in the email are the IP address of the source and destination for the session, a threat ID and the timestamp of when the event was logged.

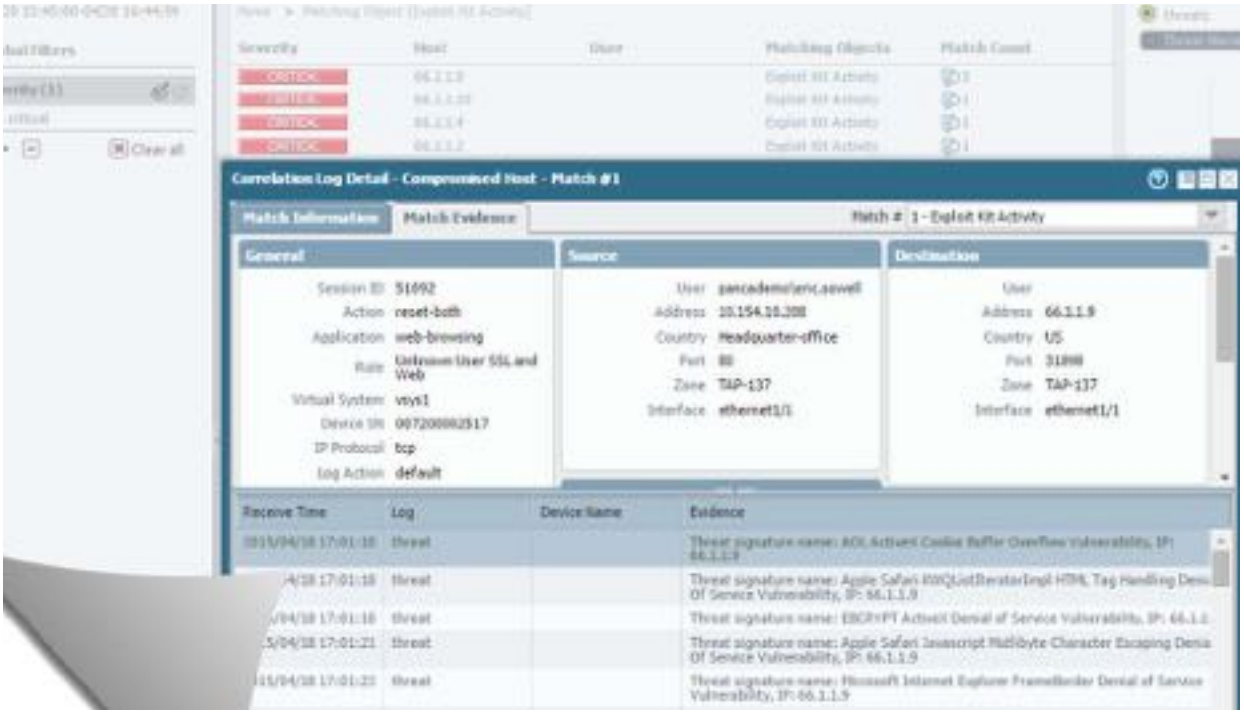
 © Palo Alto Networks, Inc.

Review the Widgets in the ACC

In the **ACC> Threat Activity** tab, check the **Compromised Hosts** widget and **Threat Activity** widget for any critical or high severity threats.



Look into the matching object and click the match count link to view the match evidence for the incident.



Review Threat Logs

To begin investigating the alert, use the threat ID to search the threat logs on Panorama (**Monitor > Logs > Threat**). From the threat logs, you can find the IP address of the victim, export the packet capture (PCAP, has a green arrow icon in the log entry) and use a network analyzer tool such as WireShark to review the packet details. In the HTTP case, look for a malformed or bogus HTTP REFERER in the protocol, suspicious host, URL strings, the user agent, the IP address and port in order to validate the incident. Data from these pcaps is also useful in searching for similar data patterns and creating custom signatures or modifying security policy to better address the threat in the future.



As a result of this manual review, if you feel confident about the signature, consider transitioning the signature from an alert action to a block action for a more aggressive approach. In some cases, you may choose to add the attacker IP to an IP block list to prevent further traffic from that IP address from reaching the internal network.

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

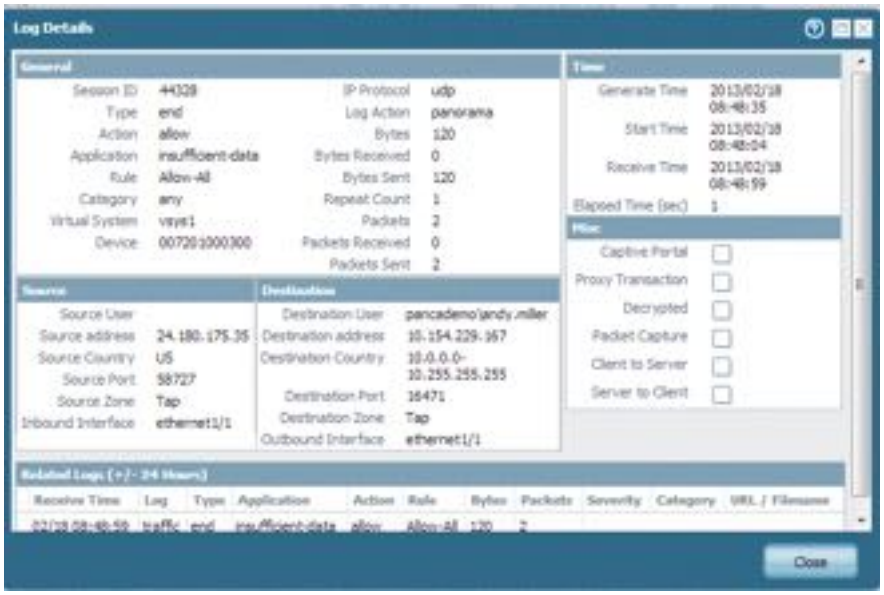
If you see this issue, enable the DNS sinkholing action in the Anti-Spyware profile in security rules to identify the infected hosts on your network. DNS sinkholing allows you to control outbound connections to malicious domains and redirect DNS queries to an internal IP address that is unused; the sinkhole that does not put out a response. When a compromised host initiates a connection to a malicious domain, instead of going out to the Internet, the firewall redirects the request to the IP address you defined and it is sinkholed. Now, reviewing the traffic logs for all hosts that connected to the sinkhole allows you locate all compromised hosts and take remedial action to prevent the spread.

To continue with the investigation on the incident, use the information on the attacker and the victim IP address to find out more information, such as:

- Where is the attacker located geographically? Is the IP address an individual IP address or a NATed IP address?
- Was the event caused by a user being tricked into going to a website, a download, or was it sent through an email attachment?
- Is the malware being propagated? Are there other compromised hosts/endpoints on the network?
- Is it a zero-day vulnerability?

The log details for each log entry display the **Related Logs** for the event. This information points you to the traffic, threat, URL filtering or other logs that you can review and correlate the events that led to the incident. For example, filter the traffic log (**Monitor > Logs > Traffic**) using the IP address as both the source and the destination IP to get a complete picture of all the external and internal hosts/clients with which this victim IP address has established a connection.

 © Palo Alto Networks, Inc.



Review WildFire Logs

In addition to the Threat logs, use the victim IP address to filter though the WildFire Submissions logs. The WildFire Submissions logs contain information on files uploaded to the WildFire service for analysis. Because spyware typically embeds itself covertly, reviewing the WildFire Submissions logs tells you whether the victim recently downloaded a suspicious file. The WildFire forensics report displays information on the URL from which the file or .exe was obtained, and the behavior of the content. It informs you if the file is malicious, if it modified registry keys, read/wrote into files, created new files, opened network communication channels, caused application crashes, spawned processes, downloaded files, or exhibited other malicious behavior. Use this information to determine whether to block the application that caused the infection (web-browsing, SMTP, FTP), make more stringent URL Filtering rules, or restrict some applications/actions (for example, file downloads to specific user groups).

Access to the WildFire logs from Panorama requires the following: a WildFire subscription, a File Blocking profile that is attached to a Security rule, and Threat log forwarding to Panorama.

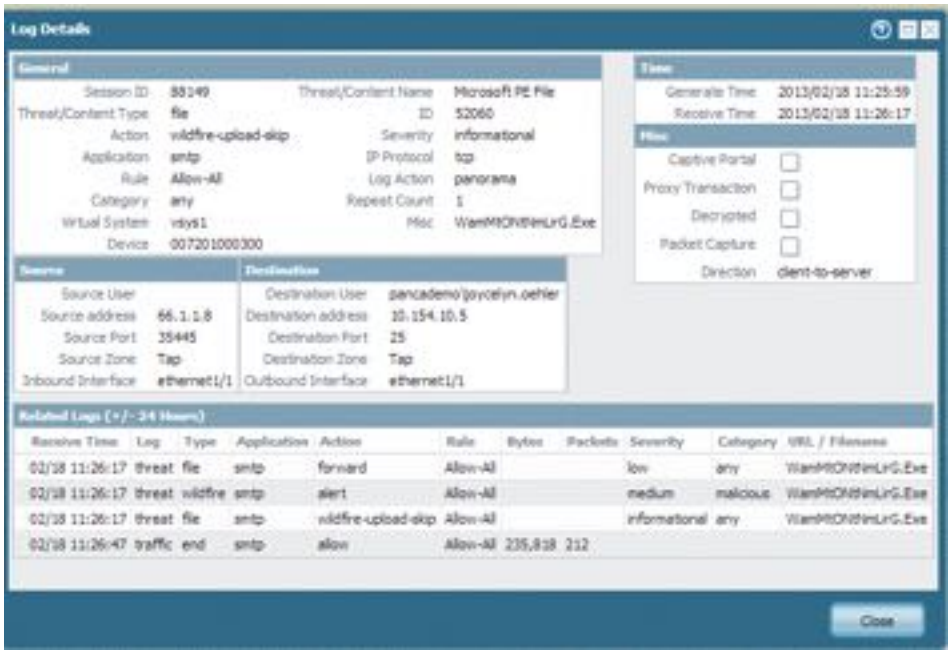
If Panorama will manage firewalls running software versions earlier than PAN-OS 7.0, specify a WildFire server from which Panorama can gather analysis information for WildFire samples that those firewalls submit. Panorama uses the information to complete WildFire Submissions logs that are missing field values introduced in PAN-OS 7.0. Firewalls running earlier releases won't populate those fields. To specify the server, select **Panorama > Setup > WildFire**, edit the General Settings, and enter the **WildFire Server** name. The default is **wildfire-public-cloud**, which is the WildFire cloud hosted in the United States.

If WildFire determines that a file is malicious, a new antivirus signature is created within 24-48 hours and made available to you. If you have a WildFire subscription, the signature is made available within 30-60 minutes as part of the next WildFire signature update. As soon as the Palo Alto Networks next-generation firewall has received a signature for it, if your configuration is configured to block malware, the file will be blocked and the information on the blocked file will be visible in your threat logs. This process is tightly integrated to protect you from this threat and stems the spread of malware on your network.

© Palo Alto Networks, Inc.

Review Data Filtering Logs

The Data Filtering log (**Monitor > Logs > Data Filtering**) is another valuable source for investigating malicious network activity. While you can periodically review the logs for all the files that you are being alerted on, you can also use the logs to trace file and data transfers to or from the victim IP address or user, and verify the direction and flow of traffic: server to client or client to server. To recreate the events that preceded and followed an event, filter the logs for the victim IP address as a destination, and review the logs for network activity.



Because Panorama aggregates information from all managed firewalls, it presents a good overview of all activity in your network. Some of the other visual tools that you can use to survey traffic on your network are the **Threat Map**, **Traffic Map**, and the **Threat Monitor**. The threat map and traffic map (**Monitor > AppScope > Threat Map** or **Traffic Map**) allow you to visualize the geographic regions for incoming and outgoing traffic. It is particularly useful for viewing unusual activity that could indicate a possible attack from outside, such as a DDoS attack. If, for example, you do not have many business transactions with Eastern Europe, and the map reveals an abnormal level of traffic to that region, click into the corresponding area of the map to launch and view the ACC information on the top applications, traffic details on the session count, bytes sent and received, top sources and destinations, users or IP addresses, and the severity of the threats detected, if any. The threat monitor (**Monitor > AppScope > Threat Monitor**) displays the top ten threats on your network, or the list of top attackers or top victims on the network.

Update Security Rules

With all the information you have now uncovered, you can sketch together how the threat impacts your network—the scale of the attack, the source, the compromised hosts, the risk factor—and evaluate what changes, if any, to follow through. Here are some suggestions to consider:

 © Palo Alto Networks, Inc.

- Forestall DDoS attacks by enhancing your DoS Protection profile to configure random early drop or to drop SYN cookies for TCP floods. Consider placing limits on ICMP and UDP traffic. Evaluate the options available to you based on the trends and patterns you noticed in your logs, and implement the changes using Panorama templates.

  Create a dynamic block list (**Objects > Dynamic Block Lists**), to block specific IP addresses that you have uncovered from several intelligence sources: analysis of your own threat logs, DDoS attacks from specific IP addresses, or a third-party IP block list.

  The list must be a text file that is located on a web server. Using device groups on Panorama, push the object to the managed firewalls so that the firewalls can access the web server and import the list at a defined frequency. After creating a dynamic block list object, define a Security rule that uses the address object in the source and destination fields to block traffic from or to the IP address, range, or subnet defined. This approach allows you to block intruders until you resolve the issue and make larger policy changes to secure your network.
- Determine whether to create shared policy rules or device group rules to block specific applications that caused the infection (web-browsing, SMTP, FTP), make more stringent URL Filtering rules, or restrict some applications/actions (for example, file downloads to specific user groups).
- On Panorama, you can also switch to the device context and configure the firewall for Botnet reports that identify potential botnet-infected hosts on the network.

 © Palo Alto Networks, Inc.

Panorama High Availability

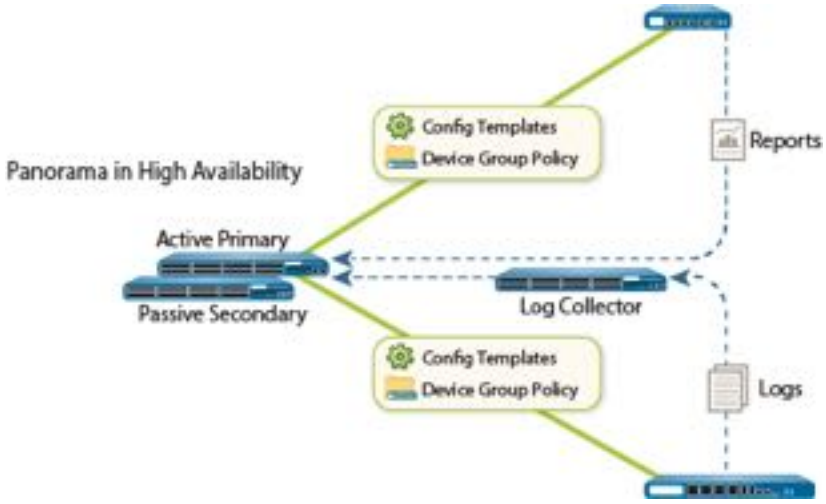
Panorama High Availability (HA) is a configuration in which two Panorama servers are placed in a group (two-device cluster) to provide redundancy in the event of a system or network failure. Panorama in HA provides continuity in the task of centrally administering and monitoring the firewalls to secure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

- Panorama HA Prerequisites
- Priority and Failover on Panorama in HA
- Failover Triggers
- Logging Considerations in Panorama HA
- Synchronization Between Panorama HA Peers
- Manage a Panorama HA Pair

Panorama HA Prerequisites

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

- **The same form factor**—Must both be hardware-based appliances (M-Series appliances) or virtual appliances. The M-Series appliances must be the same model: both M-100 appliances or both M-500 appliances. For HA, the M-Series appliances must be in Panorama mode; M-Series appliances in Log Collector mode do not support HA.
- **The same Panorama OS version**—Must run the same Panorama version to synchronize configuration information and maintain parity for a seamless failover.
- **The same set of licenses**—Must have the same device management capacity license.
- **(Panorama virtual appliance only) Unique serial number**—Must have unique serial numbers; if the serial number is the same for both Panorama instances, they will be in suspended mode until you resolve the issue.



The Panorama servers in the HA configuration are peers and you can use either (active or passive) to centrally manage the devices with a few exceptions (see Synchronization Between Panorama HA Peers). The HA peers use the management port to synchronize the configuration elements pushed to the managed devices and to maintain state information. Typically, Panorama HA peers are geographically located in different sites, so you need to make sure that the management port IP address assigned to each peer is routable through your network. HA connectivity uses TCP port 28 with encryption enabled. If encryption is not enabled, ports 28769 and 28260 are used for HA connectivity and to synchronize configuration between the HA peers. We recommend less than 500ms latency between the peers. To determine the latency, use Ping during a period of normal traffic.

 © Palo Alto Networks, Inc.

Priority and Failover on Panorama in HA

Each Panorama peer in the HA pair is assigned a *priority* value. The priority value of the primary or secondary peer determines which will be eligible for being the main point of administration and log management. The peer set as primary assumes the active state, and the secondary becomes passive. The active peer handles all the configuration changes and pushes them to the managed firewalls; the passive peer cannot make any configuration changes or push configuration to the managed firewalls. However, either peer can be used to run reports or to perform log queries.

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

When a failover occurs, only the state (active or passive) of the device changes; the priority (primary and secondary) does not. For example, when the primary peer fails, its status changes from active-primary to passive- primary.

A peer in the active-secondary state can perform all functions with two exceptions:

- It cannot manage device deployment functions such as license updates or software upgrades on the managed firewalls.
- It cannot log to an NFS until you manually change its priority to primary. (Panorama virtual appliance only)

The following table lists the capabilities of Panorama based on its state and priority settings:

Capability	active-primary	passive-primary passive-secondary	active-secondary
Switch device context	■	■	■
Perform distributed reporting	■	■	■
Manage shared policy	■	■	■
Log to local disk	■	■ (Optional on the Panorama virtual appliance only)	■ (Optional on the Panorama virtual appliance only)
Log to an NFS partition (Panorama virtual appliance only)	■	■	■
Deploy software and licenses	■	■	■
Export Panorama configuration	■	■	■

For more information, see Panorama HA Prerequisites or Set Up HA on Panorama.

 © Palo Alto Networks, Inc.

Failover Triggers

When a failure occurs on the active device and the passive device takes over the task of managing the firewalls, the event is called a failover. A failover is triggered when a monitored metric on the active device fails. This failure transitions the state on the primary Panorama from active-primary to passive-primary, and the secondary Panorama becomes active-secondary.

The conditions that trigger a failover are:

- The Panorama peers cannot communicate with each other and the active peer does not respond to health and status polls; the metric used is HA Heartbeat Polling and Hello Messages.

  When the Panorama peers cannot communicate with each other, the active peer monitors whether the devices are still connected to it before a failover is triggered. This check helps in avoiding a failover and causing a split-brain scenario, where both Panorama peers are in an active state.
- One or more of the destinations (IP addresses) specified on the active peer cannot be reached; the metric used is HA Path Monitoring.

In addition to the failover triggers listed above, a failover also occurs when the administrator places the device is a suspended state or if *preemption* occurs. Preemption is a preference for the primary Panorama to resume the active role after recovering from a failure (or user-initiated suspension). By default, preemption is enabled and when the primary Panorama recovers from a failure and becomes available, the secondary Panorama relinquishes control and returns to the passive state. When preemption occurs, the event is logged in the System log.

If you are logging to an NFS datastore, do not disable preemption because it allows the primary peer (that is mounted to the NFS) to resume the active role and write to the NFS datastore. For all other deployments, preemption is only required if you want to make sure that a specific device is the preferred active device.

HA Heartbeat Polling and Hello Messages

The HA peers use hello messages and heartbeats to verify that the peer is responsive and operational. Hello messages are sent from one peer to the other at the configured Hello Interval to verify the state of the other. The heartbeat is an ICMP ping to the HA peer, and the peer responds to the ping to establish that the devices are connected and responsive. By default, the interval for the heartbeat is 1000 milliseconds and 8000ms for hello messages.

HA Path Monitoring

Path monitoring checks the network connectivity and link state for an IP address or group of IP addresses (path group). The active peer uses ICMP pings to verify that one or more destination IP addresses can be reached. For example, you can monitor the availability of interconnected networking devices like a router or a switch, connectivity to a server, or some other vital device that is in the flow of traffic. Make sure that the node/device configured for monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a path monitoring failure and trigger a failover.

© Palo Alto Networks, Inc.

The default ping interval is 5000ms. An IP address is considered unreachable when three consecutive pings (the default value) fail, and a device failure is triggered when any or all of the IP addresses monitored become unreachable. By default, if any one of the IP addresses becomes unreachable, the HA state transitions to non-functional.

 © Palo Alto Networks, Inc.

Logging Considerations in Panorama HA

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



By default, when the managed devices cannot connect to Panorama (M-Series appliance and the Panorama virtual appliance), they buffer the logs; when the connection is restored, they resume sending logs from where it was last left off.

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

- Logging Failover on a Panorama Virtual Appliance
- Logging Failover on an M-Series Appliance

Logging Failover on a Panorama Virtual Appliance

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

Log Storage Type	Description
Virtual disk	By default, the managed devices send logs as independent streams to each Panorama HA peer. By default, if a peer becomes unavailable, the managed devices buffer the logs and when the peer reconnects it resumes sending logs from where it had left off (subject to disk storage capacity and duration of the disconnection). Logging to a virtual disk provides redundancy in logging. However, the maximum log storage capacity is 2TB.  The option to forward logs only to the active peer is configurable (see Modify Log Forwarding and Buffering Defaults). However, Panorama does not support log aggregation across the HA pair. So, if you log to a virtual disk or local disk, for monitoring and reporting you must query the Panorama peer that collects the logs from the managed devices.
Network File System (NFS)	When configured to use an NFS, only the active-primary device mounts to the NFS-based log partition and can receive logs. On failover, the primary device goes into a passive-primary state. In this scenario, until preemption occurs, the active-secondary Panorama manages the devices, but it does not receive the logs and it cannot write to the NFS. To allow the active-secondary peer to log to the NFS, you must manually switch it to primary so that it can mount to the NFS partition. For instructions, see Switch Priority after Panorama Failover to Resume NFS Logging.

Logging Failover on an M-Series Appliance

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

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

For more information, see Panorama HA Prerequisites or Set Up HA on Panorama.

 © Palo Alto Networks, Inc.

Synchronization Between Panorama HA Peers

The Panorama HA peers synchronize the running configuration each time you commit changes on the active Panorama peer. The candidate configuration is synchronized between the peers each time you save the configuration on the active peer or just before a failover occurs.

Settings that are common across the pair, such as shared objects and policy rules, device group objects and rules, template configuration, and administrative access configuration, are synchronized between the Panorama HA peers.

The settings that are not synchronized are those that are unique to each peer, such as the following:

- Panorama HA configuration—Priority setting, peer IP address, path monitoring groups and IP addresses
- Panorama configuration—Management port IP address, FQDN settings, login banner, NTP server, time zone, geographic location, DNS server, permitted IP addresses for accessing Panorama, Simple Network Management Protocol (SNMP) system settings, and dynamic content update schedules
- Scheduled configuration exports
- NFS partition configuration and all disk quota allocation for logging
- Disk quota allocation for the different types of logs and databases on the Panorama local storage (SSD)



If you use a master key to encrypt the private keys and certificates on Panorama, you must use the same master key on both HA peers. If the master keys differ, Panorama cannot synchronize the HA peers.

For more information, see Panorama HA Prerequisites or Set Up HA on Panorama.

Manage a Panorama HA Pair

- Set Up HA on Panorama
- Test Panorama HA Failover
- Switch Priority after Panorama Failover to Resume NFS Logging
- Restore the Primary Panorama to the Active State



To install software or content updates, see Install Updates for Panorama with HA Configuration.

Set Up HA on Panorama

Review the Panorama HA Prerequisites before performing the following steps:

Set Up HA on Panorama		
Step 1	Set up connectivity between the MGT ports on the HA peers.	The Panorama peers communicate with each other using the MGT port. Make sure that the IP addresses you assign to the MGT port on the Panorama servers in the HA pair are routable and that the peers can communicate with each other across your network. To set up the MGT port, see Perform Initial Configuration of the Panorama Virtual Appliance or Perform Initial Configuration of the M-Series Appliance. **Pick a device in the pair and complete the remaining tasks.**
Step 2	Enable HA and (optionally) enable encryption for the HA connection.	1. Select **Panorama > High Availability** and edit the **Setup** section. 2. Select **Enable HA.** 3. In the **Peer HA IP Address** field, enter the IP address assigned to the peer device. 4. In the **Monitor Hold Time** field, enter the length of time (milliseconds) that the system will wait before acting on a control link failure (range is 1000-60000, default is 3000). 5. If you do not want encryption, clear the **Encryption Enabled** check box and click **OK**: no more steps are required. If you do want encryption, select the **Encryption Enabled** check box, click **OK**, and perform the following tasks: a. Select **Panorama > Certificate Management > Certificates**. b. Select **Export HA key**. Save the HA key to a network location that the peer device can access. c. On the peer device, navigate to **Panorama > Certificate Management > Certificates**, select **Import HA key**, browse to the location where you saved the key, and import it.

 © Palo Alto Networks, Inc.

Set Up HA on Panorama (Continued)		
Step 3	Set the HA priority.	1. In **Panorama > High Availability**, edit the **Election Settings** section. 2. Define the **Device Priority** as **Primary** or **Secondary**. Make sure to set one peer as primary and the other as secondary. If both peers have the same priority setting, the peer with the higher serial number will be placed in a suspended state. 3. Define the **Preemptive** behavior. By default preemption is enabled. The preemption selection— enabled or disabled—must be the same on both peers. If you are using an NFS for logging and you have disabled preemption, to resume logging to the NFS see Switch Priority after Panorama Failover to Resume NFS Logging.
Step 4	To configure path monitoring, define one or more path groups. The path group lists the destination IP addresses (nodes) that Panorama must ping to verify network connectivity.	Perform the following steps for each path group that includes the nodes that you want to monitor. 1. Select **Panorama > High Availability** and, in the Path Group section, click **Add**. 2. Enter a **Name** for the path group. 3. Select a **Failure Condition** for this group: • **any** triggers a link monitoring failure if any one of the IP addresses becomes unreachable. • **all** triggers a link monitoring failure only when none of the IP addresses are reachable. 4. **Add** each destination IP address you want to monitor. 5. Click **OK**. The Path Group section displays the new group.
Step 5	(Optional) Select the failure condition for path monitoring on Panorama.	1. Select **Panorama > High Availability** and edit the Path Monitoring section. 2. Select a **Failure Condition**: • **all** triggers a failover only when all monitored path groups fail. • **any** triggers a failover when any monitored path group fails. 3. Click **OK**.
Step 6	Save your configuration changes.	Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.
Step 7	Configure the other Panorama peer.	Repeat Step 2 through Step 6 on the other peer in the HA pair.
Step 8	Verify that the Panorama servers are paired in HA.	After you configure both Panorama servers for HA: 1. Access the **Dashboard** on each Panorama, and view the **High Availability** widget. 2. Verify the Panorama servers are paired and synchronized: • Active Panorama—The state of the **Local** peer must be **active** and the **Running Config** must be **synchronized**. • Passive Panorama—The state of the **Local** peer must be **passive** and the **Running Config** must be **synchronized.**

Test Panorama HA Failover

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

Test Panorama HA Failover		
Step 1	Log in to the active Panorama peer.	You can verify the state of the Panorama server in the bottom right corner of the web interface.
Step 2	Suspend the active Panorama peer.	Select **Panorama > High Availability**, and then click the **Suspend local Panorama** link in the Operational Commands section.
Step 3	Verify that the passive Panorama peer has taken over as active.	On the Panorama **Dashboard**, **High Availability** widget, verify that the state of the **Local** passive server is **active** and the state of the **Peer** is **suspended**.
Step 4	Restore the suspended peer to a functional state. Wait for a couple minutes, and then verify that preemption has occurred, if preemptive is enabled.	On the Panorama you previously suspended: 1. Select **Panorama > High Availability** and, in the Operational Commands section, click **Make local Panorama functional**. 2. In the **High Availability** widget on the **Dashboard**, confirm that this (Local) Panorama has taken over as the active peer and that the other peer is now in a passive state.

Switch Priority after Panorama Failover to Resume NFS Logging

The Panorama virtual appliance running on an ESXi server can use a Network File System (NFS) datastore for logging. In an HA configuration, only the primary Panorama peer is mounted to the NFS-based log partition and can write to the NFS. When a failover occurs, and the passive Panorama becomes active, its state is active-secondary. Although a secondary Panorama peer can actively manage the devices, it cannot receive logs or write to the NFS because it does not own the NFS partition. When the managed device cannot forward logs to the primary Panorama peer, the logs are written to the local disk on each device. The devices maintain a pointer for the last set of log entries that were forwarded to Panorama so that when the passive-primary Panorama becomes available again, they can resume forwarding logs to it.

Use the instructions in this section to manually switch priority on the active-secondary Panorama peer so that it can begin logging to the NFS partition. The typical scenarios in which you might need to trigger this change are as follows:

- Preemption is disabled. By default, preemption is enabled on Panorama and the primary peer resumes as active when it becomes available again. When preemption is disabled, you need to switch the priority on the secondary peer to primary so that it can mount the NFS partition, receive logs from the managed devices, and write to the NFS partition.
- The active Panorama fails and cannot recover from the failure in the short term. If you do not switch the priority, when the maximum log storage capacity on the firewall is reached, the oldest logs will be overwritten to enable it to continue logging to its local disk. This situation can lead to loss of logs.

Switch Priority after Panorama Failover to Resume NFS Logging	
Step 1	Log in to the currently passive-primary Panorama, select **Panorama > Setup > Operations** and, in the Device Operations section, click **Shutdown Panorama**.

 © Palo Alto Networks, Inc.

Switch Priority after Panorama Failover to Resume NFS Logging	
Step 2	Log in to the active-secondary Panorama, select **Panorama > High Availability**, edit the Election Settings, and set the **Priority** to **Primary**.
Step 3	Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. Do not reboot when prompted.
Step 4	Log in to the Panorama CLI and enter the following command to change the ownership of the NFS partition to this peer: `request high-availability convert-to-primary`
Step 5	Select **Panorama > Setup > Operations** and, in the Device Operations section, click **Reboot Panorama**.
Step 6	Power on the Panorama peer that you powered off in Step 1. This peer will now be in a passive-secondary state.

Restore the Primary Panorama to the Active State

By default, the preemptive capability on Panorama allows the primary Panorama to resume functioning as the active peer as soon as it becomes available. However, if preemption is disabled, the only way to force the primary Panorama to become active after recovering from a failure, a non-functional, or a suspended state, is by suspending the secondary Panorama peer.

Before the active-secondary Panorama goes into a suspended state, it transfers the candidate configuration to the passive device so that all your uncommitted configuration changes are saved and can be accessed on the other peer.

Suspend the Secondary Panorama		
Step 1	Suspend Panorama.	1. Log in to the Panorama peer that you want to place in a suspended state. 2. Select **Panorama > High Availability**, and click the **Suspend local Panorama** link in the Operational Commands section.
Step 2	Verify that the status displays that the device was suspended at user request.	On the **Dashboard**, **High Availability** widget, verify that the **Local** state is **suspended**. A failover is triggered when you suspend a peer, and the other Panorama takes over as the active peer.

Restore the Primary Panorama to a Functional State		
Step 3	Restore the suspended Panorama to a functional state.	1. In the **Panorama > High Availability** tab, Operational Commands section, click the **Make local Panorama functional** link. 2. On the **Dashboard**, **High Availability** widget, confirm that the device has transitioned to either the active or passive state.

 © Palo Alto Networks, Inc.

Administer Panorama

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

- Manage Configuration Backups
- Compare Changes in Panorama Configurations
- Validate a Panorama Configuration
- Restrict Access to Configuration Changes
- Add Custom Logos to Panorama
- View Panorama Task Completion History
- Manage Storage Quotas and Expiration Periods for Logs and Reports
- Monitor Panorama
- Reboot or Shut Down Panorama
- Configure Panorama Password Profiles and Complexity



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

Manage Configuration Backups

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

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

You can perform the following tasks to manage configuration backups:

▲ Schedule Export of Configuration Files
▲ Manage Panorama and Firewall Configuration Backups
▲ Configure the Number of Configuration Backups Panorama Stores
▲ Load a Configuration Backup on a Managed Firewall

Schedule Export of Configuration Files

Panorama saves a backup of its running configuration as well as the running configurations of all managed firewalls. The backups are in XML format with file names that are based on serial numbers (of Panorama or the firewalls). Use these instructions to schedule daily exports of the backups to a remote host. Panorama exports the backups as a single gzip file. You require superuser privileges to schedule the export.



If Panorama has a high availability (HA) configuration, you must perform these instructions on each peer to ensure the scheduled exports continue after a failover. Panorama does not synchronize scheduled configuration exports between HA peers.

Schedule the Export of Configuration Files	
Step 1	Select **Panorama > Scheduled Config Export** and click **Add**.
Step 2	Enter a **Name** and **Description** for the scheduled file export and **Enable** it.
Step 3	Using the 24-hour clock format, enter a daily **Scheduled Export Start Time** or select one from the drop-down.
Step 4	Set the export **Protocol** to Secure Copy (**SCP**) or File Transfer Protocol (**FTP**).
Step 5	Enter the details for accessing the server, including: **Hostname** or IP address, **Port**, **Path** for uploading the file, **Username**, and **Password**.

 © Palo Alto Networks, Inc.

Schedule the Export of Configuration Files (Continued)	
Step 6	(**SCP** only) Click **Test SCP server connection.** To enable the secure transfer of data, you must verify and accept the host key of the SCP server. Panorama doesn't establish the connection until you accept the host key. If Panorama has an HA configuration, perform this step on each HA peer so that each one accepts the host key of the SCP server. If Panorama can successfully connect to the SCP server, it creates and uploads the test file named **ssh-export-test.txt**.
Step 7	Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

Manage Panorama and Firewall Configuration Backups

Manage Panorama and Firewall Configuration Backups	
Step 1	Select **Panorama > Setup > Operations**.

Manage Panorama and Firewall Configuration Backups	
Step 2	In the Configuration Management section, select from the following options: • **Revert to last saved Panorama configuration**—Overwrites the current candidate configuration and restores the last saved candidate configuration from disk. • **Revert to running Panorama configuration**—Reverts all changes saved to the candidate configuration; it effectively allows you to undo all configuration changes that were made since the last commit operation. • **Save named Panorama configuration snapshot**—Saves the candidate configuration to a file. Enter a file name or select an existing file to overwrite. Note that the current active configuration file (running-config.xml) cannot be overwritten. • **Save candidate Panorama configuration**—Saves the candidate configuration to disk; it is the same as using the **Save** link at the top of the page to save the changes to the candidate configuration file. • **Load Panorama configuration version**—Loads a configuration file from a list of previously committed versions. • **Load named Panorama configuration snapshot**—Loads a selected candidate configuration; you can select a previously imported or saved configuration. The current candidate configuration is overwritten. • **Export named Panorama configuration snapshot**—Exports the active configuration (running-config.xml) or a previously saved or imported configuration. Select the configuration file to be exported. You can open the file and/or save it in any network location. • **Export Panorama configuration version**—Exports a previously committed version of the configuration file. Select the version to export. • **Export Panorama and devices config bundle**—This option is used to manually generate and export the latest version of the configuration backup of Panorama and that of each managed firewall. To automate the process of creating and exporting the configuration bundle daily to a Secure Copy (SCP) or FTP server, see Schedule Export of Configuration Files. • **Export or push device config bundle**—After you import a firewall configuration into Panorama, Panorama creates a firewall configuration bundle named <firewall_name>_import.tgz, in which all local policies and objects are removed. You can then click **Export or push device config bundle** to perform one of the following actions: **Push & Commit** the configuration bundle to the firewall. This action cleans the firewall (removes any local configuration from it), enabling you to manage the firewall from Panorama. **Export** the configuration to the firewall without loading it. When you are ready to load the configuration, log in to the firewall CLI and run the configuration mode command `load device-state`. This command cleans the firewall in the same way as the **Push & Commit** option. • **Import device configuration into Panorama**—Enables you to Migrate a Firewall to Panorama Management. • **Import named Panorama configuration snapshot**—Imports a previously exported configuration file. Click **Browse** to locate the saved file and click **OK** to import.

Configure the Number of Configuration Backups Panorama Stores

Configure the Number of Configuration Backups Panorama Stores	
Step 1	Select **Panorama > Setup > Management** and edit the Logging and Reporting Settings.
Step 2	For the **Number of Versions for Config Backups**, enter a value between 1 and 1048576. The default is 100.
Step 3	Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

 © Palo Alto Networks, Inc.

Load a Configuration Backup on a Managed Firewall

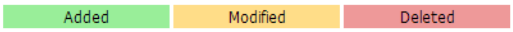
Use Panorama to load a configuration backup on a managed firewall. You can choose to revert to a previously saved or committed configuration on the firewall. Panorama pushes the selected version to the managed firewall, and the current candidate configuration on the firewall is overwritten.

Load a Configuration Backup on a Managed Firewall	
Step 1	Select **Panorama > Managed Devices**.
Step 2	Select the **Manage...** link in the **Backups** column.
Step 3	Select from the **Saved Configurations** or the **Committed Configurations**. • Click the link in the **Version** column to view the contents of the selected version. • Click **Load** to load a chosen configuration version.
Step 4	Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

Compare Changes in Panorama Configurations

To compare configuration changes on Panorama, you can select any two sets of configuration files: the candidate configuration, the running configuration, or any other configuration version that has been previously saved or committed on Panorama. The side-by-side comparison allows you to:

- Preview the changes in configuration before committing them to Panorama. You can, for example, preview the changes between the candidate configuration and the running configuration. As a best practice, select the older version on the left pane and the newer version on the right pane, to easily compare and identify modifications.
- Perform a *configuration audit* to review and compare the changes between two sets of configuration files.

Compare Changes in Panorama Configurations	
Step 1	Select **Panorama > Config Audit.**
Step 2	For each drop-down, select a configuration for the comparison.
Step 3	Select the number of lines that you want to include for **Context**, and click **Go**. To easily compare versions, the changes have color highlighting. 

Configure the Number of Versions Panorama Stores for Configuration Audits	
Step 1	Select **Panorama > Setup > Management** and edit the Logging and Reporting Settings.
Step 2	For the **Number of Versions for Config Audit**, enter a value between 1 and 1048576. The default is 100.
Step 3	Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

View and Compare Panorama Configuration Files Before Committing
Click **Commit,** select **Preview Changes**, select the number of lines of context you want to see, and click **OK**.

 © Palo Alto Networks, Inc.

Validate a Panorama Configuration

You can perform a syntactic validation (of configuration syntax) and semantic validation (whether the configuration is complete and makes sense) of a Panorama candidate configuration before committing it. The results display all the errors and warnings of a commit, including rule shadowing and application dependency warnings. For example, the validation could indicate an invalid route destination or a missing account and password that are required to query a server. This gives you the chance to fix errors before committing. You can independently validate the configurations of Panorama, device groups, and templates.

Validate a Panorama Configuration	
• Validate a Panorama candidate configuration.	1. After modifying the Panorama configuration, click **Commit**. 2. For the **Commit Type**, select **Panorama**. 3. Click **Validate Changes**. If the validation succeeds, the **Result** field displays **OK**. The **Details** field displays any errors or warnings.
• Validate a candidate configuration for a device group or template.	1. Click **Commit**, for **Commit Type**, select **Template** or **Device Group**, and select a template or device group. 2. (Optional) Select the **Merge with Device Candidate Config** check box. 3. (Optional—device groups only) Select the **Include Device and Network Templates** check box. 4. Click **Validate Changes**. The **Status** column displays any errors or warnings, and clicking the text displays details about them.

Restrict Access to Configuration Changes

Use locks to prevent multiple administrative users from making configuration changes or committing changes on Panorama, shared policy rules, or to selected templates and/or device groups.

- Types of Locks
- Locations for Taking a Lock
- Take a Lock
- View Lock Holders
- Enable Automatic Acquisition of the Commit Lock
- Remove a Lock

Types of Locks

The available lock types are:

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

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



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

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

Locations for Taking a Lock

The administrator can take a lock for any of the following:

- **Device group**—Restricts changes to the selected device group but not its descendant device groups.
- **Template**—Restricts changes to the firewalls included in the selected template. (You can't take a lock for a template stack, only for individual templates within the stack.)
- **Shared**—Restricts changes to the centrally administered rules—pre-rules and post-rules—that are shared across all device groups. For more information on shared rules, see Device Group Policies.

 © Palo Alto Networks, Inc.

- **Panorama**—Restricts access to changes on Panorama.

Take a Lock

Take a Lock	
Step 1	Click the lock icon at the top right of the web interface.
Step 2	Select **Take Lock**.
Step 3	For the **Type**, based on your role/permissions, select **Commit** or **Config**.
Step 4	As a best practice, add a **Comment** to describe the reasons for taking the lock.
Step 5	Click **OK** and **Close**.

View Lock Holders

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

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

Enable Automatic Acquisition of the Commit Lock

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

Enable Automatic Acquisition of the Commit Lock	
Step 1	Select **Panorama > Setup > Management** tab and edit the General Settings.
Step 2	Select the **Automatically Acquire Commit Lock** check box.
Step 3	Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

Remove a Lock

Remove a Lock	
Step 1	Click the lock icon at the top right of the web interface.
Step 2	Select the lock that you want to release, click **Remove Lock**, and click **OK**.  Unless you are a superuser, you can remove only the lock that you previously took.

 © Palo Alto Networks, Inc.

Add Custom Logos to Panorama

You can upload image files to customize the following areas on Panorama:

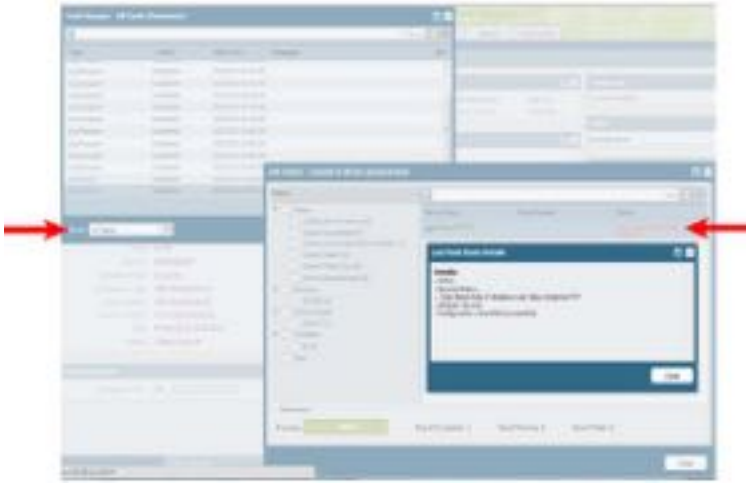
- Background image on the login screen
- Header on the top left corner of the web interface; you can also hide the Panorama default background
- Title page and footer image in PDF reports

Supported image types include .jpg, .gif, and .png. Image files for use in PDF reports cannot contain an alpha channel. The size of the image must be less than 128 Kilobytes (131,072 bytes); the recommended dimensions are displayed on screen. If the dimension is larger than the recommended size, the image will be automatically cropped.

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

View Panorama Task Completion History

Use the Task Manager to view currently-running tasks, historical task data, event success or failure information, and related errors.

View Panorama Task Completion History	
Step 1	Click the **Tasks** icon on the bottom right corner of the web interface.
Step 2	Select the list of tasks to review. By default **All Tasks** are displayed. 
Step 3	You can filter by **All** or **Running** tasks and select **Jobs**, **Reports**, or **Log Requests**: • **Jobs**—Lists commits, auto commits, downloads and installs for software and dynamic updates performed on locally on Panorama or centrally pushed to the managed firewalls from Panorama. Each job is a link; click the link in the Type column to view details on the firewalls, status, and review errors, if any. • **Reports**—Displays the status and start time for scheduled reports. • **Log Requests**—Lists the log queries triggered from the **Monitor >Log Viewer** tab or the **Dashboard**. For example, to display the logs in the URL Filtering widget or the Data Filtering widget on the Dashboard, log requests are generated on Panorama.

 © Palo Alto Networks, Inc.

Manage Storage Quotas and Expiration Periods for Logs and Reports

- Log and Report Storage
- Log and Report Expiration Periods
- Configure Storage Quotas and Expiration Periods for Logs and Reports

Log and Report Storage

You can edit the default storage quotas for each log type but not for reports. When a log quota reaches the maximum size, Panorama starts overwriting the oldest log entries with the new log entries. The Panorama virtual appliance and M-Series appliance have different locations for storing logs and different predefined storage capacities for reports:

- **Panorama virtual appliance**—Panorama writes all logs to its assigned storage space, which can be any of one the following:
  - The approximately 11GB storage allocated by default on the virtual disk that you created when installing Panorama.
  - An additional virtual disk: see Add a Virtual Disk to Panorama on an ESXi Server or Add a Virtual Disk to Panorama in vCloud Air.
  - An NFS partition: see Mount the Panorama ESXi Server to an NFS Datastore.

  The storage space for reports is 200MB.
- **M-Series appliance**—Panorama saves logs to its internal SSD and RAID-enabled disks. The M-Series appliance uses its internal SSD to store the Config logs and System logs that Panorama and its Log Collectors generate, and also to store the Application Statistics (App Stats) logs that Panorama automatically receives at 15 minute intervals from all managed firewalls. Panorama saves all other log types to its RAID-enabled disks. The RAID disks are either local to the M-Series appliance in Panorama mode or are in a Dedicated Log Collector (M-Series appliance in Log Collector mode). To edit the storage quotas for logs on the RAID disks, you must modify the Collector Group configuration. The storage space for reports is 500MB for Panorama 6.1 or later releases and 200 MB for earlier releases.

Log and Report Expiration Periods

You can configure automatic deletion based on time for the logs that the Panorama management server and Log Collectors collect from firewalls, as well as the logs and reports that Panorama and the Log Collectors generate locally. This is useful in deployments where periodically deleting monitored information is desired or necessary. For example, deleting user information after a certain period might be mandatory in your organization for legal reasons. You configure separate expiration periods for:

- Reports—Panorama deletes reports nightly at 2:00 a.m., when it generates scheduled reports.
- Each log type—Panorama evaluates logs as it receives them, and deletes logs that exceed the configured expiration period.

- Each summary log type—Panorama evaluates logs after the various summary periods (hourly, daily, and weekly), and deletes logs that exceed the configured expiration period.

> Weekly summary logs that fall short of the expiration threshold when log deletion occurs could age past the threshold before the next log deletion. For example, if you configure Traffic Summary logs to expire after 20 days and a weekly Traffic Summary log is 19 days old when the device deletes expired logs, the device doesn't delete that log. The next time the device checks for weekly logs to delete, 7 days later, that log will be 26 days old.
>
> Panorama synchronizes expiration periods across high availability (HA) pairs. Because only the active HA peer generates logs, the passive peer has no logs or reports to delete unless failover occurs and it starts generating logs.
>
> Even if you don't set expiration periods, when a log quota reaches the maximum size, Panorama starts overwriting the oldest log entries with the new log entries.

Configure Storage Quotas and Expiration Periods for Logs and Reports

Configure Storage Quotas and Expiration Periods for Logs and Reports	
Step 1 Configure the storage quotas and expiration periods for: • Logs of all types that a Panorama virtual appliance receives from firewalls. • App Stats logs that Panorama (a virtual appliance or M-Series appliance) receives from firewalls. • System and Config logs that Panorama (a virtual appliance or M-Series appliance) and its Log Collectors generate locally. The Panorama management server stores these logs. If you reduce a storage quota such that the current logs exceed it, after you commit the change, Panorama removes the oldest logs to fit the quota.	1. Select **Panorama > Setup > Management** and edit the Logging and Reporting Settings. 2. In the **Log Storage** tab, enter the storage **Quota (%)** for each log type. When you change a percentage value, the page refreshes to display the corresponding absolute value (Quota GB/MB column) based on the total allotted storage on Panorama. 3. Enter the **Max Days** (expiration period) for each log type (range is 1-2,000). By default, the fields are blank, which means the logs never expire. To reset the quotas and expiration periods to the factory defaults, click **Restore Quota Defaults** at the bottom right of the dialog.
Step 2 Configure the expiration period for reports that Panorama (a virtual appliance or M-Series appliance) generates.	1. Select the **Log Export and Reporting** tab. 2. Enter the **Report Expiration Period** in days (range is 1-2,000). By default, the field is blank, which means reports never expire. 3. Click **OK** to save your changes.

 © Palo Alto Networks, Inc.

Configure Storage Quotas and Expiration Periods for Logs and Reports (Continued)		
Step 3	Configure the storage quotas and expiration periods for logs of all types (except App Stats logs) that a Panorama M-Series appliance receives from firewalls. The Log Collectors store these logs. You configure these storage quotas at the Collector Group level, not for individual Log Collectors.	1. Select **Panorama > Collector Groups** and select the Collector Group. 2. In the **General** tab, click the **Log Storage** value. This field doesn't display a value unless you assigned Log Collectors to the Collector Group. If the field displays 0MB after you assign Log Collectors, verify that you enabled the disk pairs when configuring the Log Collector and that you committed the changes (**Panorama > Managed Collectors** > **Disks**). 3. Enter the storage **Quota(%)** for each log type. When you change a percentage value, the page refreshes to display the corresponding absolute value (Quota GB/MB column) based on the total storage allotted to the Collector Group. 4. Enter the **Max Days** (expiration period) for each log type (range is 1-2,000). By default, the fields are blank, which means the logs never expire. To reset the quotas and expiration periods to the factory defaults, click **Restore Quota Defaults** at the bottom right of the dialog. 5. Click **OK** to save your changes.
Step 4	Commit your changes.	1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 2. (M-Series appliance only) Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Group you modified, and click **OK**.
Step 5	Verify that Panorama applied the storage quota changes.	1. Select **Panorama > Setup > Management** and, in the Logging and Reporting Settings, verify that the **Log Storage** values are correct for the logs that the Panorama management server stores. 2. Select **Panorama > Collector Groups**, select the Collector Group you modified, and verify that the **Log Storage** values in the **General** tab are correct for the logs that the Log Collectors store. You can also verify the Collector Group storage quotas by logging in to a Log Collector CLI and entering the operational command `show log-diskquota-pct`.

Monitor Panorama

To monitor Panorama and its managed collectors, you can periodically view their System and Config logs (filter the log data by device), configure a Simple Network Management Protocol (SNMP) manager to collect (GET) Panorama statistics on a regular basis, or configure SNMP traps or email alerts that notify you when a monitored metric changes state or reaches a threshold on Panorama. Email alerts and SNMP traps are useful for immediate notification about critical system events that need your attention. To configure email alerts or SNMP traps, see Configure Log Forwarding from Panorama to External Destinations.

- Panorama System and Configuration Logs
- Monitor Panorama and Log Collector Statistics Using SNMP

Panorama System and Configuration Logs

You can configure Panorama to send notifications when a system event or configuration change occurs. By default, Panorama logs every configuration change to the Config logs. In the System logs, each event has a severity level to indicate its urgency and impact. When you Configure Log Forwarding from Panorama to External Destinations, you can forward all system events or just events of certain severity levels. The following table summarizes the severity levels:

Severity	Description
Critical	Indicates a failure and the need for immediate attention, such as a hardware failure, including high availability (HA) failover and link failures.
High	Serious issues that will impair the operation of the system, including disconnection of a Log Collector or a commit failure.
Medium	Mid-level notifications, such as Antivirus package upgrades, or a Collector Group commit.
Low	Minor severity notifications, such as user password changes.
Informational	Notification events such as log in or log out, any configuration change, authentication success and failure notifications, commit success, and all other events that the other severity levels don't cover.

The M-Series appliance stores Config and System logs on its SSD. The Panorama virtual appliance stores the logs on the assigned storage volume (see Set Up the Panorama Virtual Appliance). If you need longer-term log storage for auditing, you can also Configure Log Forwarding from Panorama to External Destinations.



For information on using Panorama to monitor firewall logs, see Monitor Network Activity.

 © Palo Alto Networks, Inc.

Monitor Panorama and Log Collector Statistics Using SNMP

You can configure an SNMP manager to request statistics from a Panorama management server and configure Panorama to respond. For example, the SNMP manager can request the high availability (HA) mode, Panorama state, and Panorama version. If the Panorama management server is an M-Series appliance in Panorama mode (not a virtual appliance), it can also provide logging statistics: average logs per second, storage duration of each log type, and log disk usage. Panorama doesn't synchronize SNMP configurations between HA peers; you must enable SNMP requests and responses on each peer.

You can also configure a Dedicated Log Collector (M-Series appliance in Log Collector mode) to respond to requests for statistics such as connection status, disk drive metrics, software version, average CPU, average logs per second, and log storage duration for each log type. This information is useful when evaluating whether you need to expand log storage capacity.



You can't configure an SNMP manager to control Panorama or Log Collectors (using SET messages); an SNMP manager can only collect statistics (using GET messages).

For details on how Palo Alto Networks devices implement SNMP, see SNMP for Palo Alto Networks Devices.

Monitor Panorama and Log Collector Statistics Using SNMP		
Step 1	Configure the SNMP Manager to get statistics from Panorama and the Log Collectors.	The following steps are an overview of the tasks you perform on the SNMP manager. For the specific steps, refer to the documentation of your SNMP manager. 1. To enable the SNMP manager to interpret device statistics, load the Supported MIBs for Palo Alto Networks devices and, if necessary, compile them. 2. For each Panorama M-Series or virtual appliance that the SNMP manager will monitor, define the connection settings (IP address and port) and authentication settings (SNMPv2c community string or SNMPv3 username and password) for the device. All Palo Alto Networks devices use port 161. The SNMP manager can use the same or different connection and authentication settings for multiple devices. The settings must match those you define when you configure SNMP on the device (see Step 4). For example, if you use SNMPv2c, the community string you define when configuring the device must match the community string you define in the SNMP manager for that device. 3. Determine the object identifiers (OIDs) of the statistics you will monitor. For example, to monitor the logging rate, a MIB browser shows that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.3.30.1.1 in PAN-PRODUCT-MIB.my. For details, see Use an SNMP Manager to Explore MIBs and Objects. 4. Configure the SNMP manager to monitor the desired OIDs.
Step 2	Enable SNMP traffic on the management (MGT) interface of the Panorama management server.	1. Select **Panorama > Setup > Management** and edit the Management Interface Settings. 2. In the Services section, select the **SNMP** check box and click **OK**.

Monitor Panorama and Log Collector Statistics Using SNMP (Continued)		
Step 3	Enable SNMP traffic on the management (MGT) interface of any M-Series appliances in Log Collector mode:	1. Select **Panorama > Managed Collectors** and select the Log Collector. 2. Select the **Management** tab, select the **SNMP** check box, and click **OK**.
Step 4	Configure the Panorama management server to respond to statistics requests from an SNMP manager.	1. Select **Panorama > Setup > Operations** and, in the Miscellaneous section, click **SNMP Setup**. 2. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP for Palo Alto Networks Devices. • **V2c**—Enter the **SNMP Community String**, which identifies a community of SNMP managers and monitored devices (Panorama, in this case), and serves as a password to authenticate the community members to each other. Don't use the default community string `public`; it is well known and therefore not secure. • **V3**—Create at least one SNMP view group and one user. User accounts and views provide authentication, privacy, and access control when SNMP managers get device statistics. – **Views**—Each view is a paired OID and bitwise mask: the OID specifies a MIB, and the mask (in hexadecimal format) specifies which objects are accessible inside (include matching) or outside (exclude matching) that MIB. Click **Add** in the first list and enter a **Name** for the group of views. For each view in the group, click **Add** and configure the view **Name**, **OID**, matching **Option** (**include** or **exclude**), and **Mask**. – **Users**: Click **Add** in the second list, enter a username in the Users column, select the **View** group from the drop-down, enter the authentication password (**Auth Password**) used to authenticate to the SNMP manager, and enter the privacy password (**Priv Password**) used to encrypt SNMP messages to the SNMP manager. 3. Click **OK** to save the settings.
Step 5	Configure the Dedicated Log Collectors (if any) to respond to SNMP requests.	For each Collector Group: 1. Select **Panorama > Collector Groups** and select the Collector Group. 2. Select the **Monitoring** tab, configure the same settings as in Step 4, and click **OK**.
Step 6	Commit your changes.	1. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 2. Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Groups you edited, and click **Commit** again.
Step 7	Monitor the Panorama and Log Collector statistics in an SNMP manager.	Refer to the documentation of your SNMP manager.

 © Palo Alto Networks, Inc.

Reboot or Shut Down Panorama

The reboot option initiates a graceful restart of Panorama. A shutdown halts the system and powers it off. To restart Panorama, after a shutdown, manually disconnect and re-cable the power cord on the system.

Reboot or Shut Down Panorama	
Step 1	Select **Panorama > Setup > Operations**.
Step 2	In the Device Operations section, select **Reboot Panorama** or **Shutdown Panorama**.

Configure Panorama Password Profiles and Complexity

To secure the local administrator account, you can define password complexity requirements that are enforced when administrators change or create new passwords. Unlike password profiles, which can be applied to individual accounts, the password complexity rules are firewall-wide and apply to all passwords.

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

Configure Panorama Password Profiles and Complexity	
Step 1 Configure minimum password complexity settings.	1. Select **Panorama > Setup > Management** and edit the Minimum Password Complexity section. 2. Select **Enabled**. 3. Define the **Password Format Requirements**. You can enforce the requirements for uppercase, lowercase, numeric, and special characters that a password must contain. 4. To prevent the account username (or reversed version of the name) from being used in the password, select **Block Username Inclusion (including reversed)**.  5. Define the password **Functionality Requirements**. If you have configured a password profile for an administrator, the values defined in the password profile will override the values that you have defined in this section. 

 © Palo Alto Networks, Inc.

Configure Panorama Password Profiles and Complexity (Continued)		
Step 2	Create password profiles. You can create multiple password profiles and apply them to administrator accounts as required to enforce security.	1. Select **Panorama > Password Profiles** and click **Add**. 2. Enter a **Name** for the password profile and define the following: a. **Required Password Change Period**: Frequency, in days, at which the passwords must be changed. b. **Expiration Warning Period**: Number of days before expiration that the administrator will receive a password reminder. c. **Post Expiration Grace Period:** Number of days that the administrator can still log in to the system after the password expires. d. **Post Expiration Admin Login Count**: Number of times that the administrator can log in to the system after the password has expired.

 © Palo Alto Networks, Inc.

Troubleshooting

The following topics address Panorama issues:

- Troubleshoot Panorama System Issues
- Troubleshoot Log Storage and Connection Issues
- Replace an RMA Firewall
- Troubleshoot Commit Failures
- Troubleshoot Registration or Serial Number Errors
- Troubleshoot Reporting Errors
- View Task Success or Failure Status

Troubleshoot Panorama System Issues

- Generate Diagnostic Files for Panorama
- Diagnose Panorama Suspended State
- Monitor the File System Integrity Check
- Manage Panorama Storage for Software and Content Updates
- Recover from Split Brain in Panorama HA Deployments

Generate Diagnostic Files for Panorama

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

Generate Diagnostic Files for Panorama	
Step 1	Select **Panorama > Support** and click **Generate Tech Support File**.
Step 2	Download and save the file to your computer.
Step 3	Upload the file to your case on the Palo Alto Networks Support Portal.

Diagnose Panorama Suspended State

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

- Verify that the serial number on each Panorama virtual appliance is unique. If the same serial number is used to create two or more instances of Panorama, all instances using the same serial number will be suspended.
- Verify that you have set the HA priority setting on one peer as *Primary* and the other as *Secondary*. If the priority setting is identical on both peers, the Panorama peer with a higher numerical value in serial number is placed in a suspended state.
- Verify that both Panorama HA peers are running the same Panorama version (major and minor version number).

Monitor the File System Integrity Check

Panorama periodically performs a file system integrity check (FSCK) to prevent corruption of the Panorama system files. This check occurs after eight reboots or at a reboot that occurs 90 days after the last FSCK was executed. If Panorama is running a FSCK, the web interface and Secure Shell (SSH) login screens will display

 © Palo Alto Networks, Inc.

a warning to indicate that an FSCK is in progress. You cannot log in until this process completes. The time to complete this process varies by the size of the storage system; depending on the size, it can take several hours before you can log back in to Panorama.

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

Manage Panorama Storage for Software and Content Updates

You can Install Content and Software Updates for Panorama and Deploy Updates to Firewalls and Log Collectors Using Panorama. You cannot configure the amount of space available on Panorama to store updates. When the alloted storage capacity reaches 90%, Panorama alerts you to free up space (delete stored updates) for new downloads or uploads. The maximum number of updates is a global setting that applies to all the updates that Panorama stores. You can use only the CLI to configure this setting. The default value is two updates of each type.

Manage Panorama Storage for Software and Content Updates	
• Modify the maximum number of updates of each type.	Access the Panorama CLI and enter the following, where `<number>` can be between 2 and 64: `> set max-num-images count <number>`
• View the number of updates that Panorama currently stores.	Enter: > **show max-num-images**
• Use the web interface to delete updates to free up space on Panorama.	1. Select the type of update to delete: • Firewall or Log Collector updates: – PAN-OS/Panorama software images—Select **Panorama > Device Deployment > Software**. – GlobalProtect agent/app software updates—Select **Panorama > Device Deployment > GlobalProtect Client**. – Content updates—Select **Panorama > Device Deployment > Dynamic Updates**. • Panorama software images—Select **Panorama > Software**. • Panorama content updates—Select **Panorama > Dynamic Updates**. 2. Click the **X** icon in the far right column for the image or update.

• Use the CLI to delete updates to free up space on Panorama.	Delete software images by version:
	`> delete software version <version_number>`
	Delete content updates:
	`> delete content update <filename>`

Recover from Split Brain in Panorama HA Deployments

When Panorama is configured in a high availability (HA) setup, the managed firewalls are connected to both the active and passive Panorama HA peers. When the connection between the active and the passive Panorama peers fails, before the passive Panorama takes over as the active peer it checks whether any firewall is connected to both the active and the passive peer. If even one firewall is connected to both peers, the failover is not triggered.

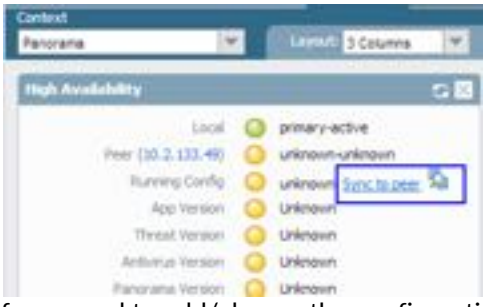
In the rare event that a failover is triggered when a set of firewalls are connected to the active peer and a set of firewalls are connected to the passive peer, but none of the firewalls are connected to both peers, it is called a split brain. When a split brain occurs, the following conditions occur:

- Neither Panorama peer is aware of the state nor the HA role of the other peer.
- Both Panorama peers become active and manage a unique set of firewalls.

To resolve a split brain, debug your network issues and restore connectivity between the Panorama HA peers.

However, if you need to make configuration changes to your firewalls without restoring the connection between the peers, here are a couple of options:

- Manually add the same configuration changes on both Panorama peers. This ensures that when the link is reestablished the configuration is synchronized.
- If you need to add/change the configuration at only one Panorama location, make the changes and sync the configuration (make sure that you initiate the sync from the peer on which you made the changes) when the link between the Panorama peers is re-established.



- If you need to add/change the configuration for only the connected firewalls at each location, you can make configuration changes independently on each Panorama peer. Because the peers are disconnected, there is no replication and each peer now has a completely different configuration file (they are out of sync). Therefore, to ensure that the configuration changes on each peer are not lost when the connection is restored, you cannot allow the configuration to be automatically re-synchronized. To solve this problem, export the configuration from each Panorama peer and manually merge the changes using an external diff and merge tool. After the changes are integrated, you can import the unified configuration file on the primary Panorama and then synchronize the imported configuration file with the peer.

 © Palo Alto Networks, Inc.

Troubleshoot Log Storage and Connection Issues

- Verify Panorama Port Usage
- Resolve Zero Log Storage for a Collector Group
- Replace a Failed Disk on an M-Series Appliance
- Replace the Virtual Disk on a Panorama ESXi Server
- Replace the Virtual Disk on vCloud Air
- Migrate Logs to a New M-Series Appliance in Log Collector Mode
- Migrate Logs to a New M-Series Appliance in Panorama Mode
- Recover Logs after Panorama Failure/RMA in Non-HA Deployments
- Regenerate Metadata for M-Series Appliance RAID Pairs

Verify Panorama Port Usage

To ensure that Panorama can communicate with managed firewalls, Log Collectors, and its high availability (HA) peer, use the following table to verify the ports that you must open on your network.



On an M-Series appliance running Panorama 6.1 or later releases, you can optionally assign the log collection and Collector Group communication functions to the Eth1 or Eth2 interfaces (instead of to the default MGT interface). The ports listed in the following table apply regardless of which function you assign to which interface. For example, if you assign log collection to MGT and assign Collector Group communication to Eth2, then MGT will use port 3978 and Eth2 will use port 28270. (The Panorama virtual appliance can only use the MGT interface for all these functions.)

Communicating Devices & Direction of Connection Establishment	Ports Used in Panorama 5.x	Ports Used in Panorama 6.x and later	Description
Panorama and Panorama (HA) Direction: Each peer initiates its own connection to the other	28	28	For HA connectivity and synchronization if encryption is enabled.
Panorama and Panorama (HA) Direction: Each peer initiates its own connection to the other	28769 and 28260 (5.1) 28769 and 49160 (5.0)	28260 and 28769	For HA connectivity and synchronization if encryption is not enabled.
Panorama and managed firewalls Direction: Initiated by the firewall	3978	3978	A bi-directional connection where the logs are forwarded from the firewall to Panorama; and configuration changes are pushed from Panorama to the managed firewalls. Context switching commands are sent over the same connection.

Communicating Devices & Direction of Connection Establishment	Ports Used in Panorama 5.x	Ports Used in Panorama 6.x and later	Description
Panorama and Log Collector Direction: Initiated by the Log Collector	3978	3978	For management and log collection/reporting. Used for communication between the default Log Collector on a Panorama in Panorama mode, and for communicating with Log Collectors in a distributed log collection deployment.
Log Collector to Log Collector Direction: Each Log Collector initiates a connection to the other Log Collectors in the Collector Group	49190	28270	For distributing blocks and all binary data between Log Collectors.

Resolve Zero Log Storage for a Collector Group

The log storage capacity for the Collector Group might display as 0MB if the disk pairs are not enabled for logging. You must select the Log Collector and enable the disk pairs for logging in the **Panorama > Managed Collectors** tab; for instructions, see Step 9 in the Manage Collector Groups topic.

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

Replace a Failed Disk on an M-Series Appliance

If a disk fails on the M-Series appliance, you must replace the disk and reconfigure it in a RAID pair. This allows the data to be mirrored and synchronized between the disks in the RAID pair.

Replace a Failed Disk		
Step 1	Install the new disk in the appropriate drive bay.	Refer to the M-100 or M-500 Hardware Reference Guide for instructions to replace the failed with the new disk.

Step 2	Set up the disk in a RAID pair.
	The time required to mirror the data on the drive may vary from several minutes to a couple hours, depending on the amount of data on the drive.
		This example uses the drives in the disk bays B1.
		1. Enter the following commands to add the disk to the RAID pair and confirm the request when prompted:
		`request system raid add B1`
		2. To monitor the progress of the RAID configuration and verify that the disk is RAID enabled, enter the following command:
		`show system raid detail`

 © Palo Alto Networks, Inc.

Replace the Virtual Disk on a Panorama ESXi Server

You can't resize a virtual disk after adding it to a Panorama virtual appliance on an ESXi server. Because the Panorama virtual appliance allows only one log storage location, if you need to increase or decrease disk space for logging, you must replace the virtual disk on the ESXi server to adjust the log storage capacity.



You will lose the logs on the existing disk when you replace it. One way to preserve the logs is to set up a new Panorama virtual appliance for the new disk (see Set Up the Panorama Virtual Appliance) and maintain access to the Panorama containing the old disk for as long as you need its logs. A second way to preserve the logs is to copy them from the old disk to the new disk. Copying can take several hours, depending on how many logs the disk currently stores, and Panorama cannot collect logs during the process. Contact Palo Alto Networks Customer Support for instructions on how to copy logs between disks. A third way to preserve existing logs is to Configure Log Forwarding from Panorama to External Destinations before you replace the virtual disk.

Replace the Virtual Disk on a Panorama ESXi Server		
Step 1	Remove the old virtual disk.	1. Access the VMware vSphere Client and select the **Virtual Machines** tab. 2. Right-click the Panorama virtual appliance and select **Power > Power Off**. 3. Right-click the Panorama virtual appliance and select **Edit Settings**. 4. Select the virtual disk in the **Hardware** tab and click **Remove**. 5. Select one of the Removal Options and click **OK**.
Step 2	Add the new virtual disk.	1. Add a Virtual Disk to Panorama on an ESXi Server. Panorama running on ESXi 5.5 and later versions supports a virtual disk of up to 8TB. Panorama running on an earlier ESXi version supports a virtual disk of up to 2TB. 2. In the vSphere Client, right-click the Panorama virtual appliance and select **Power > Power On**. The reboot process might take several minutes and the message `cache data unavailable` will display.
Step 3	Verify that the modified log storage capacity is correct.	1. Log in to the Panorama virtual appliance. 2. Select **Panorama > Setup > Management** and verify that the Logging and Reporting Settings section, Log Storage field, displays the modified log storage capacity accurately.

Replace the Virtual Disk on vCloud Air

You can't resize a virtual disk after adding it to a Panorama virtual appliance running on VMware vCloud Air. Because the Panorama virtual appliance allows only one log storage location, you must replace the virtual disk to modify the log storage capacity.

You will lose the logs on the existing disk when you replace it. One way to preserve the logs is to set up a new Panorama virtual appliance for the new disk (see Set Up the Panorama Virtual Appliance) and maintain access to the Panorama containing the old disk for as long as you need its logs. A second way to preserve the logs is to copy them from the old disk to the new disk. Copying can take several hours, depending on how many logs the disk currently stores, and Panorama cannot collect logs during the process. Contact Palo Alto Networks Customer Support for instructions on how to copy logs between disks. A third way to preserve existing logs is to Configure Log Forwarding from Panorama to External Destinations before you replace the virtual disk.

Replace the Virtual Disk on vCloud Air	
Step 1 Remove the old virtual disk.	1. Access the vCloud Air web console and select your **Virtual Private Cloud OnDemand** region. 2. Select the Panorama virtual appliance in the **Virtual Machines** tab. 3. Select **Actions > Edit Resources**. 4. Click **x** for the virtual disk you are removing.
Step 2 Add the new virtual disk.	1. **Add another disk**. 2. Set the **Storage** to up to 8TB and set the storage tier to **Standard** or **SSD-Accelerated**. 3. **Save** your changes.
Step 3 Reboot Panorama.	1. Log in to the Panorama virtual appliance. 2. Select **Panorama > Setup > Operations** and **Reboot Panorama**.
Step 4 Verify that the modified log storage capacity is correct.	1. Log in to the Panorama virtual appliance after it reboots. 2. Select **Panorama > Setup > Management** and verify that the Logging and Reporting Settings section, Log Storage field, displays the modified log storage capacity accurately.

Migrate Logs to a New M-Series Appliance in Log Collector Mode

If you need to replace an M-Series appliance in Log Collector mode (Dedicated Log Collector), you can migrate the logs it collected from firewalls by moving its RAID disks to a new M-Series appliance. This enables you to recover logs after a system failure on the M-Series appliance or to migrate logs as part of a hardware upgrade (from an M-100 appliance to an M-500 appliance). This procedure applies whether the Panorama management server that manages the Dedicated Log Collector is a Panorama virtual appliance or an M-Series appliance in Panorama mode.

 © Palo Alto Networks, Inc.

Migrate Logs to a New M-Series Appliance in Log Collector Mode	

Step 1 Perform initial setup of the new M-Series appliance that will be a Dedicated Log Collector.	1. Rack mount the M-Series appliance. Refer to the M-100 or M-500 Appliance Hardware Reference Guide for instructions.
	2. Perform Initial Configuration of the M-Series Appliance.
	When configuring interfaces, configure only the Management (MGT) interface. Switching to Log Collector mode removes any Eth1 and Eth2 interface configurations. If the Log Collector will use Eth1 and Eth2, add them when configuring the Log Collector in the next step.
	3. Register Panorama.
	4. Transfer licenses as follows only if the new M-Series appliance is the same hardware model as the old M-Series appliance. Otherwise, you must purchase new licenses.
	a. Log in to the Customer Support Portal.
	b. Select the **Assets** tab and click the **Spares** link.
	c. Click the Serial Number of the new M-Series appliance.
	d. Click **Transfer Licenses**.
	e. **Select** the old M-Series appliance and click **Submit**.
	5. Activate a Panorama Support License.
	6. Activate a device management license. If you are migrating from an M-100 appliance to an M-500 appliance, enter the auth-code associated with the migration license.
	7. Install Content and Software Updates for Panorama.
	The M-500 appliance requires Panorama 7.0 or a later release. For important details about software versions, see Panorama, Log Collector, and Firewall Version Compatibility.
	8. Switch from Panorama mode to Log Collector mode:
	a. Access the Log Collector CLI and switch to Log Collector mode:
	`> request system system-mode logger`
	b. Enter `Y` to confirm the mode change. The M-Series appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the M-Series appliance to display the Panorama login prompt.
	If you see a `CMS Login` prompt, press Enter without typing a username or password.
	9. Use the Log Collector CLI to enable connectivity between the Log Collector and Panorama management server. <IPaddress1 is for the MGT interface of the primary Panorama and <IPaddress2> is for the MGT interface of the secondary Panorama.
	`> configure`
	`# set deviceconfig system panorama-server <IPaddress1> panorama-server-2 <IPaddress2>`
	`# commit`
	`# exit`

Migrate Logs to a New M-Series Appliance in Log Collector Mode (Continued)	
Step 2 On the Panorama management server, add the new Log Collector as a managed collector. For all steps with commands that require a device serial number, you must type the entire serial number; pressing the Tab key won't complete a partial serial number.	1. Configure the Log Collector as a managed collector using the Panorama web interface or using the following CLI commands: > **configure** # **set log-collector <LC_serial_number> deviceconfig system hostname <LC_hostname>** # **exit** If the old Log Collector used the Eth1 and Eth2 interfaces for log collection and Collector Group communication, you must define those interfaces on the new Log Collector when you configure it as a managed collector (**Panorama > Managed Collectors > Eth1** and **Eth2**). 2. Verify that the Log Collector is connected to Panorama and that the status of its disk pairs is present/available. > **show log-collector serial-number <log-collector_SN>** The disk pairs will display as disabled at this stage of the restoration process. 3. Commit your changes to Panorama. Don't commit the changes to the Collector Group just yet. # **configure** # **commit** # **exit**
Step 3 Remove the RAID disks from the old Log Collector.	1. Power off the old Log Collector by pressing the Power button until the system shuts down. 2. Remove the disk pairs. For details, refer to the disk replacement procedure in the M-100 or M-500 Appliance Hardware Reference Guide.

 © Palo Alto Networks, Inc.

Migrate Logs to a New M-Series Appliance in Log Collector Mode (Continued)	
Step 4 Prepare the disks for migration. Generating the metadata for each disk pair rebuilds the indexes. Therefore, depending on the data size, this process can take a long time to complete. To expedite the process, you can launch multiple CLI sessions and run the metadata regeneration command in each session to complete the process simultaneously for every pair. For details, see Regenerate Metadata for M-Series Appliance RAID Pairs.	1. Insert the disks into the new Log Collector. For details, refer to the disk replacement procedure in the M-100 or M-500 Appliance Hardware Reference Guide. The disk carriers of the M-100 appliance are incompatible with those of the M-500 appliance. Therefore, when migrating between these hardware models, you must unscrew each disk from its old carrier and insert the disk in the new carrier before inserting the disk in the new appliance. You must maintain the disk pair association. Although you can place a disk pair from slot A1/A2 on the old appliance into slot B1/B2 on the new appliance, you must keep the disks together in the same slot; otherwise, Panorama might not restore the data successfully. 2. Enable the disk pairs by running the following CLI command for each pair: > **request system raid add <slot> force no-format** For example: > **request system raid add A1 force no-format** > **request system raid add A2 force no-format** The `force` and `no-format` arguments are required. The `force` argument associates the disk pair with the new Log Collector. The `no-format` argument prevents reformatting of the drives and retains the logs stored on the disks. 3. Generate the metadata for each disk pair. > **request metadata-regenerate slot <slot_number>** For example: > **request metadata-regenerate slot 1**
Step 5 Migrate the logs. You must use the Panorama CLI for this step, not the web interface. You must assign the new Log Collector to the Collector Group that contains the old Log Collector.	1. Assign the new Log Collector to the Collector Group and commit your changes to Panorama. > **configure** # **set log-collector-group <collector_group_name> logfwd-setting collectors** <new_LC_serial_number> # **commit** # **exit** 2. For each disk pair, migrate the logs from the old Log Collector to the new Log Collector and attach the disk pair to the new Log Collector. > **request log-migration from** <old_LC_serial_number> **old-disk-pair <log_disk_pair> to** <new_LC_serial_number> **new-disk-pair <log_disk_pair>** For example: > **request log-migration from 003001000010 old-disk-pair A to 00300100038 new-disk-pair A**

Migrate Logs to a New M-Series Appliance in Log Collector Mode (Continued)	
Step 6 Reconfigure the Collector Group.	1. Use the web interface to assign the new Log Collector to the firewalls that forward logs (**Panorama > Collector Groups > Device Log Forwarding**). Give the new Log Collector the same priority in the firewall preference lists as the old Log Collector. You use the web interface to perform this step because no CLI command can change the priority assignments of firewall preference lists. 2. Delete the old Log Collector from the Collector Group. > **configure** # **delete log-collector-group <group_name> logfwd-setting collectors <old_LC_serial_number>** For example: **delete log-collector-group DC-Collector-Group logfwd-setting collectors 003001000010** 3. Delete the old Log Collector from the Panorama configuration and commit your changes to Panorama. # **delete log-collector <old_LC_serial_number>** # **commit** # **exit** 4. Commit the Collector Group changes so that the managed firewalls can send logs to the new Log Collector. > **commit-all log-collector-config log-collector-group <collector_group_name>** For example: > **commit-all log-collector-config log-collector-group DC-Collector-Group**

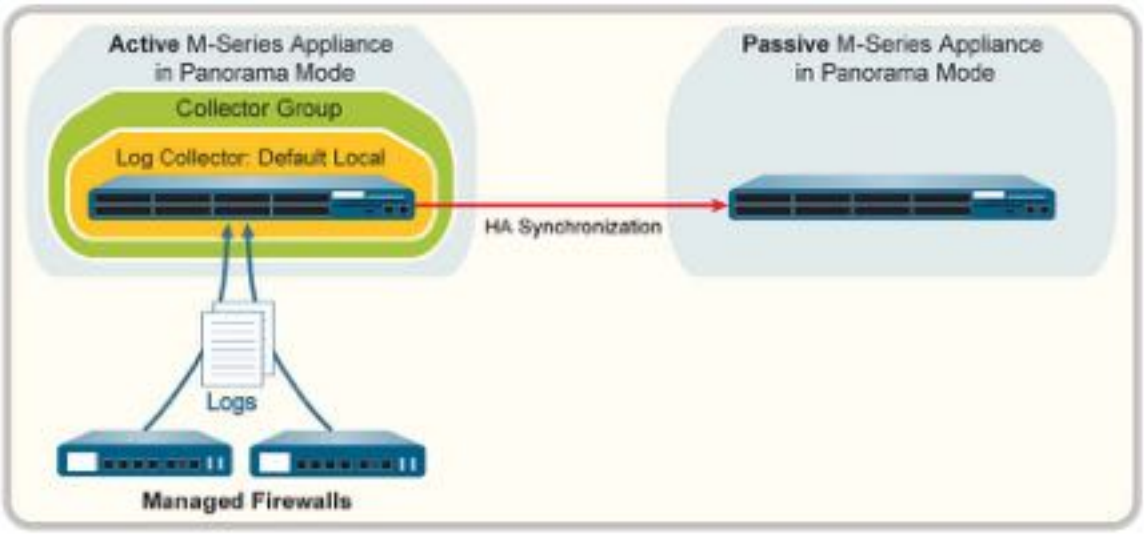
Migrate Logs to a New M-Series Appliance in Panorama Mode

If you need to replace an M-Series appliance in Panorama mode (Panorama management server), you can migrate the logs it collected from firewalls by moving its RAID disks to a new M-Series appliance, but only if Panorama is deployed in a high availability (HA) configuration. Moving the disks enables you to recover logs after a system failure on the M-Series appliance or to migrate logs as part of a hardware upgrade (from an M-100 appliance to an M-500 appliance).

This migration procedure covers the following scenarios:

- One Panorama HA peer has a default managed collector (Log Collector) in a Collector Group.

 © Palo Alto Networks, Inc.



- Both Panorama HA peers have managed collectors that belong to a single Collector Group. For details, see Figure: Multiple Default Log Collectors Per Collector Group.
- Both Panorama HA peers have a managed collector and each is assigned to a separate Collector Group. For details, see Figure: Single Default Log Collector Per Collector Group.

Migrate Logs to a New M-Series Appliance in Panorama Mode		
Step 1	Forward any logs on the SSD of the old M-Series appliance to an external destination if you want to preserve them. The SSD stores only the System and Config logs that Panorama and Log Collectors generate. You cannot move the SSD between M-Series appliances.	Configure Log Forwarding from Panorama to External Destinations.
Step 2	Remove the RAID disks from the old M-Series appliance.	1. Power off the old M-Series appliance by pressing the Power button until the system shuts down. 2. Remove the disk pairs. For details, refer to the disk replacement procedure in the M-100 or M-500 Appliance Hardware Reference Guide.

Migrate Logs to a New M-Series Appliance in Panorama Mode (Continued)	
Step 3 Perform initial setup of the new M-Series appliance.	1. Rack mount the M-Series appliance. Refer to the M-100 or M-500 Appliance Hardware Reference Guide for instructions. 2. Perform Initial Configuration of the M-Series Appliance. If the old M-Series appliance used the Eth1 and Eth2 interfaces for log collection and Collector Group communication, you must define those interfaces during initial configuration of the new M-Series appliance (**Panorama > Setup > Management**). 3. Register Panorama. 4. Transfer licenses as follows only if the new M-Series appliance is the same hardware model as the old M-Series appliance. Otherwise, you must purchase new licenses. a. Log in to the Customer Support Portal. b. Select the **Assets** tab and click the **Spares** link. c. Click the Serial Number of the new M-Series appliance. d. Click **Transfer Licenses**. e. **Select** the old M-Series appliance and click **Submit**. 5. Activate a Panorama Support License. 6. Activate a device management license. If you are migrating from an M-100 appliance to an M-500 appliance, enter the auth-code associated with the migration license. 7. Install Content and Software Updates for Panorama. The M-500 appliance requires Panorama 7.0 or a later release. For important details about software versions, see Panorama, Log Collector, and Firewall Version Compatibility. 8. Set the HA priority. The new M-Series appliance must have the same priority as the HA peer you are replacing.

 © Palo Alto Networks, Inc.

Migrate Logs to a New M-Series Appliance in Panorama Mode (Continued)	
Step 4 Prepare the disks for migration.  Generating the metadata for each disk pair rebuilds the indexes. Therefore, depending on the data size, this process can take a long time to complete. To expedite the process, you can launch multiple CLI sessions and run the metadata regeneration command in each session to complete the process simultaneously for every pair. For details, see Regenerate Metadata for M-Series Appliance RAID Pairs.	1. Insert the disks into the new M-Series appliance. For details, refer to the disk replacement procedure in the M-100 or M-500 Appliance Hardware Reference Guide.  The disk carriers of the M-100 appliance are incompatible with those of the M-500 appliance. Therefore, when migrating between these hardware models, you must unscrew each disk from its old carrier and insert the disk in the new carrier before inserting the disk in the new appliance. You must maintain the disk pair association. Although you can place a disk pair from slot A1/A2 on the old appliance into slot B1/B2 on the new appliance, you must keep the disks together in the same slot; otherwise, Panorama might not restore the data successfully. 2. Enable the disk pairs by running the following CLI command for each pair: > **request system raid add <slot> force no-format** For example: > **request system raid add A1 force no-format** > **request system raid add A2 force no-format** The `force` and `no-format` arguments are required. The `force` argument associates the disk pair with the new appliance. The `no-format` argument prevents reformatting of the drives and retains the logs stored on the disks. 3. Generate the metadata for each disk pair. > **request metadata-regenerate slot <slot_number>** For example: > **request metadata-regenerate slot 1** 4. Synchronize the configuration of the M-Series appliance HA peers. > **request high-availability sync-to-remote running-config**
Step 5 Configure the local Log Collector on the new M-Series appliance.  For all steps with commands that require a device serial number, you must type the entire serial number; pressing the Tab key won't complete a partial serial number. Don't enable the disks on the new M-Series appliance at this point. When you successfully migrate the logs, Panorama automatically enables the disks.	1. Configure the local Log Collector as a managed collector using the Panorama web interface or using the following CLI commands: > **configure** # **set log-collector <log-collector_SN> deviceconfig system hostname <log-collector-hostname>** # **exit** 2. Verify that the local Log Collector is connected to Panorama and that the status of its disk pairs is present/available. > **show log-collector serial-number <log-collector_SN>**  The disk pairs will display as disabled at this stage of the restoration process. 3. Commit your changes to Panorama. Don't commit the changes to the Collector Group just yet. > **configure** # **commit**

Migrate Logs to a New M-Series Appliance in Panorama Mode (Continued)	
Step 6 Migrate the logs. You must use the Panorama CLI for this step, not the web interface. You must assign the local Log Collector of the new M-Series appliance to the Collector Group that contains the local Log Collector of the old M-Series appliance.	1. Add the new local Log Collector as a member of the Collector Group and commit your changes to Panorama. # **set log-collector-group <collector_group_name> logfwd-setting collectors <SN_managed_collector>** # **commit** # **exit** The old local Log Collector still appears in the list of members, because you didn't yet delete it from the configuration. 2. For each disk pair, migrate the logs to the new appliance. > **request log-migration from <old_LC_serial_number> old-disk-pair <log_disk_pair> to <new_LC_serial_number> new-disk-pair <log_disk_pair>** For example: > **request log-migration from 003001000010 old-disk-pair A to 00300100038 new-disk-pair A** 3. Commit the changes to Panorama. > **configure** # **commit**
Step 7 Reconfigure the Collector Group.	1. Use the web interface to assign the new Log Collector to the firewalls that forward logs (**Panorama > Collector Groups > Device Log Forwarding**). Give the new Log Collector the same priority in the firewall preference lists as the old Log Collector. You use the web interface to perform this step because no CLI command can change the priority assignments of firewall preference lists. 2. Delete the old Log Collector from the Collector Group. # **delete log-collector-group <group_name> logfwd-setting collectors <old_LC_serial_number>** For example: # **delete log-collector-group DC-Collector-Group logfwd-setting collectors 003001000010** 3. Delete the old Log Collector from the Panorama configuration and commit your changes to Panorama. # **delete log-collector <old_LC_serial_number>** # **commit** # **exit** 4. Commit the Collector Group changes so that the managed firewalls can send logs to the new Log Collector. > **commit-all log-collector-config log-collector-group <collector_group_name>** For example: > **commit-all log-collector-config log-collector-group DC-Collector-Group**

 © Palo Alto Networks, Inc.

Recover Logs after Panorama Failure/RMA in Non-HA Deployments

If a system failure occurs on a Panorama management server that is not deployed in a high availability (HA) configuration, use this procedure to restore the configuration on the replacement Panorama and restore access to the logs on the Dedicated Log Collectors that it manages.

Panorama maintains a ring file that maps the segments and partitions that Dedicated Log Collectors use to store logs. An M-Series appliance in Panorama mode stores the ring file on its internal SSD; a Panorama virtual appliance stores the ring file on its internal disk. When a system failure occurs, a non-HA Panorama cannot automatically recover the ring file. Therefore, when you replace Panorama, you must restore the ring file to access the logs on the Dedicated Log Collectors.



Palo Alto Networks recommends deploying Panorama in an HA configuration. The active Panorama peer automatically synchronizes the ring file to the passive peer in an HA configuration, thereby maintaining access to logs on the Dedicated Log Collectors.

Recover Logs after Panorama Failure/RMA in Non-HA Deployments	
Step 1 Perform initial setup of the new Panorama platform.	1. Rack mount the M-Series appliance if that is the new platform. Refer to the M-100 or M-500 Appliance Hardware Reference Guide for instructions. 2. Perform Initial Configuration of the M-Series Appliance or Perform Initial Configuration of the Panorama Virtual Appliance.  If the old M-Series appliance used the Eth1 and Eth2 interfaces for log collection and Collector Group communication, you must define those interfaces during initial configuration of the new M-Series appliance (**Panorama > Setup > Management**). The Panorama virtual appliance does not use Eth1 or Eth2. 3. Register Panorama. 4. Transfer licenses: a. Log in to the Customer Support Portal. b. Select the **Assets** tab and click the **Spares** link. c. Click the Serial Number of the new M-Series appliance. d. Click **Transfer Licenses**. e. **Select** the old platform and click **Submit**. 5. Activate a Panorama Support License. 6. Activate a device management license. 7. Install Content and Software Updates for Panorama.  The M-500 appliance requires Panorama 7.0 or a later release. For important details about software versions, see Panorama, Log Collector, and Firewall Version Compatibility.

Recover Logs after Panorama Failure/RMA in Non-HA Deployments (Continued)		
Step 2	Restore the configuration from the old Panorama to the replacement Panorama. This task assumes that, before the system failure occurred, you followed the recommendation to back up and export your Panorama configuration.	1. Log in to the new Panorama and select **Panorama > Setup > Operations.** 2. Click **Import named Panorama configuration snapshot**, **Browse** to the saved file, and click **OK**. 3. Click **Load named Panorama configuration snapshot** and select the file you just imported. 4. Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again. 5. Select **Panorama > Managed Collectors** and verify that the Connected column displays a check mark for the Dedicated Log Collector. If the Dedicated Log Collector doesn't appear, you must reconfigure it and its Collector Group as described in Step 4.

Step 3	Fetch the ring file to restore access to the logs stored on the Dedicated Log Collector.	1. Access the CLI of the new Panorama.
		2. Fetch the ring file:
		`> request fetch ring from log-collector <serial-number>`
		For example:
		`> request fetch ring from log-collector 009201000343`
		If you don't know the serial number of the Dedicated Log Collector, log in to its CLI and enter the `show system info` operational command.
		3. Commit your changes to the Collector Group.
		`> commit-all log-collector-config log-collector-group <collector-group-name>`

 © Palo Alto Networks, Inc.

Recover Logs after Panorama Failure/RMA in Non-HA Deployments (Continued)	

Step 4 Reconfigure the Dedicated Log Collector and Collector Group if they are missing on Panorama.	1. Access the CLI of the Dedicated Log Collector and enter the following commands to display the name of its Collector Group.
	a. Enter the command:
	`> request fetch ring from log-collector <serial_number>`
	The following error will display:
	`Server error: Failed to fetch ring info from <serial_number>`
	b. Enter the command:
	`> less mp-log ms.log`
	The following error will display:
	`Dec04 11:07:08 Error: pan_cms_convert_resp_ring_to_file(pan_ops_cms.c:3719): Current configuration does not contain group CA-Collector-Group`
	In this example, the error message indicates that the missing Collector Group has the name CA-Collector-Group.
	2. Configure the Collector Group and assign the Dedicated Log Collector to it.
	`> configure`
	`# set log-collector-group <collector-group-name>`
	`# set log-collector-group <collector-group-name> logfwd-setting collector <serial-number>`
	3. Commit the changes to Panorama but not to the Collector Group.
	`# commit`
	`# exit`
	4. Fetch the ring file from the Dedicated Log Collector:
	`> request fetch ring from log-collector <serial_number>`
	5. Commit the changes to the Collector Group:
	`> commit-all log-collector-config log-collector-group <collector-group-name>`

Regenerate Metadata for M-Series Appliance RAID Pairs

When a system failure occurs on the M-Series appliance and you need to physically move the disks from one appliance to another, regenerating the metadata is necessary. The metadata is required to locate logs on the disk; when a user issues a log query, the query consults this metadata to access the requested log data.

For each configured RAID disk pair in the M-Series appliance, you must access the appliance CLI and run the following command to regenerate the metadata:

```
request metadata-regenerate slot <slot_number>
```

For example:

```
request metadata-regenerate slot 1
```

The size of the RAID disks determines how long metadata regeneration takes. On average, it takes an hour for every 100GB. When you run the command, the CLI session is locked until the command is fully executed. You can use multiple CLI sessions to save time. For example, to replace four RAID pairs with a total of 4TB of log data, launch four CLI sessions and run the command in each session to regenerate metadata simultaneously for all the pairs/slots in about 10 hours.

During metadata regeneration, the Collector Group to which these disks belong is not available and the disk pair is not available for any logging or reporting operations (writes/queries). However, you can perform other tasks such as handling new firewall connections or managing configuration changes on the managed firewalls. All other Collector Groups that Panorama manages and that aren't part of this RMA process can perform the assigned logging and reporting functionality as normal.

 © Palo Alto Networks, Inc.

Replace an RMA Firewall

To minimize the effort required to restore the configuration on a managed firewall involving a Return Merchandise Authorization (RMA), replace the serial number of the old firewall with that of the new/replacement firewall on Panorama. To then restore the configuration on the replacement firewall, either import a firewall state that you previously generated and exported from the firewall or use Panorama to generate a *partial device state* for managed firewalls running PAN-OS 5.0 and later versions. By replacing the serial number and importing the device state, you can resume using Panorama to manage the firewall.

▲ Partial Device State Generation for Firewalls

▲ Before Starting RMA Firewall Replacement

▲ Restore the Firewall Configuration after Replacement

Partial Device State Generation for Firewalls

When you use Panorama to generate a partial device state, it replicates the configuration of the managed firewalls with a few exceptions for Large Scale VPN (LSVPN) setups. You create the partial device state by combining two facets of the configuration on a managed firewall:

- Centralized configuration managed by Panorama—Panorama maintains a snapshot of the shared policy rules and templates that it pushes to firewalls.
- Local configuration on the firewall—When a configuration change is committed, each firewall sends a copy of its local configuration file to Panorama. Panorama stores this file and uses it to compile the partial device state bundle.



In an LSVPN setup, the partial device state bundle that you generate on Panorama is not the same as the version that you export from a firewall (by selecting **Device > Setup > Operations** and clicking **Export device state**). If you manually ran the device state export or scheduled an XML API script to export the file to a remote server, you can use the exported device state in your firewall replacement workflow.

If you did not export the device state, the device state that you generate in the replacement workflow will not include the dynamic configuration information, such as the certificate details and registered firewalls, that is required to restore the complete configuration of a firewall functioning as an LSVPN portal. See Before Starting RMA Firewall Replacement for more information.

Panorama does not store the device state; you generate it on request using the CLI commands listed in Restore the Firewall Configuration after Replacement.

Before Starting RMA Firewall Replacement

- The managed firewall (that was replaced) must have been on PAN-OS 5.0.4 and later version. Panorama cannot generate the *device state* for firewalls running older PAN-OS versions.
- Record the following details about the old firewall:

  - **Serial number**—You must enter the serial number on the Support portal to transfer the licenses from the old firewall to your replacement firewall. You will also enter this information on Panorama, to replace all references to the older serial number with the serial number of the replacement firewall.
  - (Recommended) **PAN-OS version and the content database version**—Installing the same software and content database versions, including the URL database vendor allows you to create the same state on the replacement firewall. If you decide to install the latest version of the content database, you may notice differences because of updates and additions to the database. To verify the versions installed on the firewall, access the firewall system logs stored on Panorama.
- Prepare the replacement firewall for deployment. Before you import the device state bundle and restore the configuration, you must:
  - Verify that the replacement firewall is of the same model and is enabled for similar operational capability. Consider the following operational features: does it need to be enabled for multi-virtual systems, support jumbo frames, or be enabled to operate in CC or FIPS mode?
  - Configure network access, transfer the licenses, and install the appropriate PAN-OS version and the content database version.
- You must use the Panorama CLI to complete this firewall replacement process. This CLI-based workflow is available for the *superuser* and *panorama-admin* user roles.
- If you have an LSVPN configuration, and are replacing a Palo Alto Networks firewall deployed as a satellite device or as an LSVPN portal, the dynamic configuration information that is required to restore LSVPN connectivity will not be available when you restore the partial device state generated on Panorama. If you have been following the recommendation to frequently generate and export the device state for firewalls in an LSVPN configuration, use the device state that you have previously exported from the firewall itself instead of generating one on Panorama.

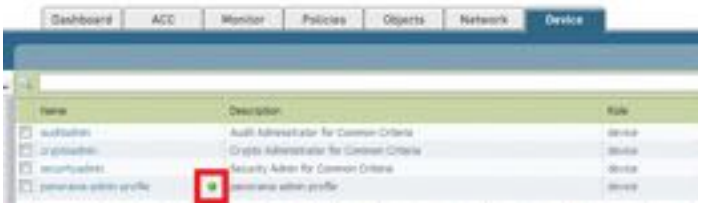
  If you have not manually exported the device state from the firewall, and need to generate a partial device state on Panorama, the missing dynamic configuration impacts the firewall replacement process as follows:
  - **If the firewall you are replacing is a portal device** that is explicitly configured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

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

 © Palo Alto Networks, Inc.

Restore the Firewall Configuration after Replacement

Restore the Firewall Configuration after Replacement	
Tasks on the new firewall:	Use the CLI for a more streamlined workflow.
Step 1 Perform initial configuration and verify network connectivity.	Use a serial port connection or a Secure Shell (SSH) connection to add an IP address, a DNS server IP address, and to verify that the firewall can access the Palo Alto Networks updates server.
Step 2 (Optional) Set the Operational mode to match that on the old firewall. A serial port connection is required for this task.	1. Enter the following CLI command to access maintenance mode on the firewall: > **debug system maintenance-mode** 2. Set the Operational mode to **Set FIPS Mode** or **Set CCEAL 4 Mode** from the main menu.
Step 3 Retrieve the license(s).	Enter the following command to retrieve your licenses: > **request license fetch**
Step 4 (Optional) Match the operational state of the new firewall with that of the old firewall. For example, enable multi-virtual system (multi-vsys) capability for a firewall that was enabled for multi-vsys capability.	Enter the commands that pertain to your firewall settings: > **set system setting multi-vsys on** > **set system setting jumbo-frame on**
Step 5 Upgrade the PAN-OS version on the firewall. You must upgrade to the same OS and content database version that is installed on the old firewall.	Enter the following commands: 1. To upgrade the content database version: > **request content upgrade download <xxx-xxxx>** 2. To install the content database version that you downloaded: > **request content upgrade install version <xxx-xxxx>** 3. To upgrade the PAN-OS software version: > **request system software download version 5.x.x** 4. To install the PAN-OS software version that you downloaded: > **request system software install version 5.x.x**
Tasks on the Panorama CLI:	You cannot perform these tasks on the Panorama web interface.
(Skip this step if you have manually exported the device state from your firewall.) Step 6 Export the device state bundle to a computer using Secure Copy (SCP) or TFTP. The export command generates the device state bundle as a tar zipped file and exports it to the specified location. This device state will not include the LSVPN dynamic configuration (satellite information and certificate details).	Enter one of the following commands: > **scp export device-state device <old serial#> to <login> @ <serverIP>: <path>** or > **tftp export device-state device <old serial#> to <login> @ <serverIP>: <path>**

Restore the Firewall Configuration after Replacement (Continued)	
Step 7 Replace the serial number of the old firewall with that of the new replacement firewall on Panorama. By replacing the serial number on Panorama you allow the new firewall to connect to Panorama after you restore the configuration on the firewall.	1. Enter the following command in Operational mode: > **replace device old <old SN#> new <new SN#>** 2. Enter Configuration mode and commit your changes. > **configure** # **commit** 3. Exit Configuration mode. # **exit**
Tasks on the new firewall:	You can use the firewall web interface to perform these tasks.
Step 8 Import the device state and commit the changes on the firewall.	1. Access the web interface of the firewall. 2. Select **Device > Setup > Operations** and click the **Import Device State** link in the Configuration Management section. 3. Browse to locate the file and click **OK**. 4. Click **Commit** to save you changes to the running configuration on the firewall. 5. To confirm that the device state restored includes the references to Panorama pushed polices and objects, verify that a little green icon appears beside the device name. 
Tasks on Panorama:	You can now use the Panorama web interface to access and manage the replaced firewall.
Step 9 Verify that you successfully restored the firewall configuration.	1. Access the Panorama web interface and select **Panorama > Managed Devices**. 2. Verify that the Connected column for the new firewall has a check mark.
Step 10 Synchronize the firewall with Panorama.	1. Click **Commit**, set the **Commit Type** to **Device Group**, select the device group that contains the firewall, select the **Include Device and Network Template** check box, and click **Commit** again. 2. (M-Series appliance only) If your firewalls forward logs to Log Collectors, click **Commit**, set the **Commit Type** to **Collector Group**, select the Collector Group that contains the firewall, and click **Commit** again. If you need to generate reports for a period when the old firewall was still functional after you installed the new firewall, you must generate a separate query for each firewall serial number because replacing the serial number on Panorama does not overwrite the information in logs.

 © Palo Alto Networks, Inc.

Troubleshoot Commit Failures

If commit failures occur on Panorama, check for the following conditions:

Symptom	Condition	Resolution
Template or device group commit failure	The ability to receive template and device groups configuration changes from Panorama is disabled on the firewall.	Access the firewall web interface, select **Device > Setup**, edit the Panorama Settings, and then click **Enable Device and Network Template** and **Enable Panorama Policy and Objects**.
Panorama, template, device group, or Collector Group commit failure	The Panorama management server has an earlier software version than the Dedicated Log Collectors or firewalls that it manages.	Upgrade the Panorama management server to the same or a higher software version than the managed devices. For details, see Panorama, Log Collector, and Firewall Version Compatibility.

Troubleshoot Registration or Serial Number Errors

On the M-Series appliance, if the **Panorama > Support** page doesn't display support license details or the **Panorama > Setup > Management** page displays Unknown for the **Serial Number** even after you Register Panorama, perform the following steps:

Resolve Registration or Serial Number Errors	
Step 1	Record the Panorama serial number from the order fulfillment email that Palo Alto Networks sent when you placed your order for Panorama.
Step 2	Select **Panorama > Setup > Management** and edit the General Settings.
Step 3	Enter the **Serial Number** and click **OK**.
Step 4	Click **Commit**, for the **Commit Type** select **Panorama**, and click **Commit** again.

 © Palo Alto Networks, Inc.

Troubleshoot Reporting Errors

If Panorama fails to generate a report, or the report is missing expected data, its content versions (for example, the Applications and Threats database) might differ from those on the managed collectors and firewalls. The content versions on Panorama must be the same as or higher than the content versions on the managed collectors and firewalls. For details, see Panorama, Log Collector, and Firewall Version Compatibility.

View Task Success or Failure Status

Use the Task Manager icon  at the bottom right of the Panorama web interface to view the success or failure of a task. The Task Manager also displays a detailed message to help debug an issue. For details, see View Panorama Task Completion History.

 © Palo Alto Networks, Inc.

PANORAMA

Security deployments are complex and overload IT teams with convoluted security rules and mountains of data from multiple sources. Panorama™ network security management empowers you with easy-to-implement, consolidated policy creation and centralized management features. Set up and control firewalls centrally with industry-leading functionality and an efficient rule base, and gain insight into network-wide traffic and threats.

Key Security Features:

MANAGEMENT:

- Deploy corporate policies centrally to be used in conjunction with local policies for maximum flexibility.
- Delegate appropriate levels of administrative control at the device level or globally with role-based management.
- Group devices into logical, hierarchical device groups for greater management flexibility.
- Utilize template stacks for easy device and network configuration.
- Easily import existing device configurations into Panorama.

VISIBILITY AND SECURITY:

- Automatically correlate indicators of threats for improved visibility and confirmation of compromised hosts across your network.
- Centrally analyze, investigate and report on network traffic, security incidents and administrative modifications.
- View a highly customizable graphical summary of applications, users, content, and security threats.
- Generate actionable, customizable reports to view application and threat traffic, SaaS usage, and user behavior across your configuration.

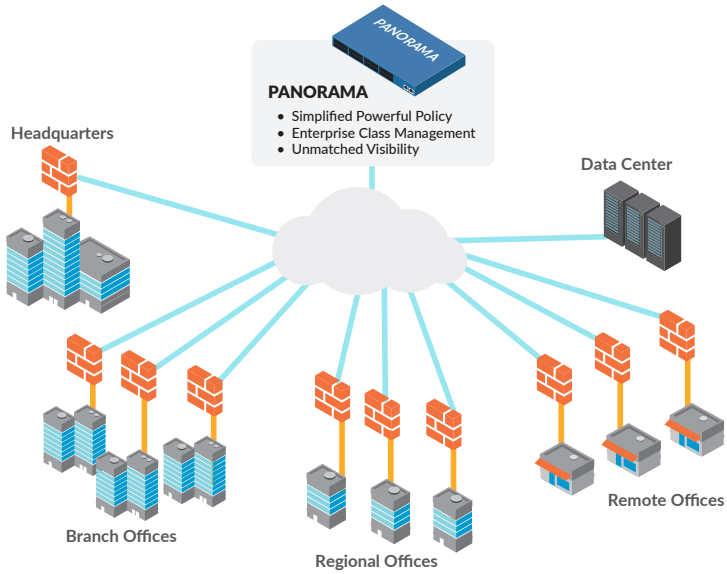


Figure 1: Panorama deployment

Simplified Powerful Policy — Panorama provides static rules in an ever-changing network and threat landscape. Manage your network security with a single security rule base for firewall, threat prevention, URL filtering, application awareness, user identification, sandboxing, file blocking, and data filtering. This crucial simplification, along with dynamic security updates, reduces workload on administrators while improving your overall security posture.

Enterprise Class Management — Panorama keeps the enterprise user in mind. Control your Internet and data center edge, and private and public cloud deployments, all from one single console. Panorama can be deployed via virtual appliances, a choice of appropriately sized appliances, or a combination of the two. Use appliances as Panorama management units, or as log collectors in hierarchical deployment options. As your network grows, you just need to add the log collectors — we take care of the rest!

Unmatched Automated Visibility and Awareness — Automated threat correlation, with a predefined set of correlation objects, cuts through the clutter of monstrous amounts of data. It identifies compromised hosts and surfaces correlated malicious behavior that would otherwise be buried in the noise of too much information. This reduces the dwell time of critical threats in your network. A clean and fully customizable Application Command Center (ACC) provides comprehensive insight into current and historical network and threat data.



Figure 1: ACC

Powerful Network Visibility: Application Command Center

Using Application Command Center (ACC) from Panorama provides you with a highly interactive, graphical view of application, URL, threat and data (files and patterns) traversing your Palo Alto Networks® firewalls. ACC includes a tabbed view of network activity, threat activity, and blocked activity, and each tab includes pertinent widgets for better visualization of traffic patterns on your network. Custom tabs can be created, which include widgets that enable you to drill down into the information that is most important to the administrator. ACC provides a comprehensive, fully customizable view of not only current but also historical data.

Additional data on URL categories and threats provides a complete and well-rounded picture of network activity. The visibility from ACC enables you to make informed policy decisions and respond quickly to potential security threats.

Reduced Response Times: Automated Correlation Engine

The automated correlation engine built into the NGFW surfaces critical threats that may be hidden in your network. It includes correlation objects that are defined by the Palo Alto Networks threat research team. These objects identify suspicious traffic patterns or a sequence of events that indicates a malicious outcome. Some correlation objects can identify dynamic patterns that have been observed from malware samples in WildFire™.

Simple Policy Control: Safely Enable Applications

Safely enabling applications means allowing access to specific applications, but protecting them with specific threat prevention, QoS, and file, data, or URL filtering policies. Panorama empowers you to set policy with a single security rule base, and simplifies the process of importing, duplicating or modifying rules across your network. The combination of centralized and local administrative control over policies and objects lets you strike a balance between consistent security at the global level and flexibility at the local level.

Enterprise Class Management

Deploying hierarchical device groups ensures that lower-level groups inherit the settings of higher-level groups. This streamlines central management and enables you to organize devices based on function and location without redundant configuration. Template stacking allows for streamlined configuration of networks and devices. Furthermore, a common user interface for both next-generation firewalls and management makes management intuitive. Features such as Global Find and Tag-based Rule Grouping empower IT administrators to take advantage of all the information in your network with ease.

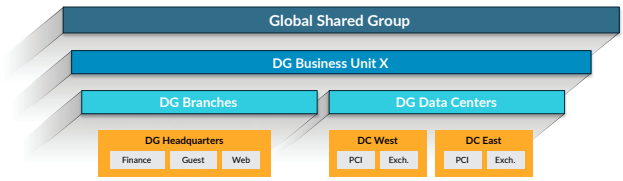


Figure 2: Device Group Hierarchy

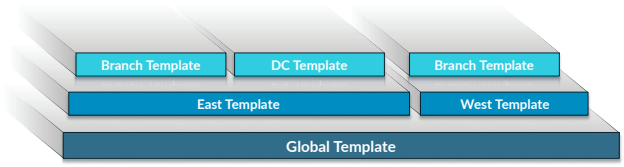


Figure 3: Template Stacking

Traffic Monitoring: Analysis, Reporting and Forensics

Panorama utilizes the same set of powerful monitoring and reporting tools available at the local device management level. As you perform log queries and generate reports, Panorama dynamically pulls the most current data directly from firewalls under management or from logs forwarded to Panorama.

- **Log Viewer:** For either an individual device or all devices, you can quickly view log activities using dynamic log filtering by clicking on a cell value and/or using the expression builder to define the sort criteria. Results can be saved for future queries or exported for further analysis.
- **Custom Reporting:** Predefined reports can be used as is, customized, or grouped together as one report in order to suit specific requirements.

- **User Activity Reports:** A user activity report shows the applications used, URL categories visited, websites visited, and all URLs visited over a specified period of time for individual users. Panorama builds the reports using an aggregate view of users' activity, no matter which firewall they are protected by, or which IP or device they may be using.
- **SaaS Reports:** A SaaS usage and threat report provides detailed visibility into all SaaS activity and related threats.
- **Log Forwarding:** Panorama aggregates logs collected from all of your Palo Alto Networks firewalls, both physical and virtual form factor, and forwards them to a remote destination for purposes such as long-term storage, forensics or compliance reporting. Panorama can forward all or selected logs, SNMP traps, and email notifications to a remote logging destination, such as a syslog server (over UDP, TCP or SSL).

Panorama Management Architecture

Panorama enables organizations to manage their Palo Alto Networks firewalls using a model that provides both central oversight and local control. Panorama provides a number of tools for centralized administration:

- **Templates/Template Stacks:** Panorama manages common device and network configuration through templates. Templates can be used to manage configuration centrally and then push the changes to all managed firewalls. This approach avoids making the same individual firewall change repeatedly across many devices. To make things even easier, templates can be stacked and used like building blocks during device and network configuration.
- **Hierarchical Device Groups:** Panorama manages common policies and objects through hierarchical device groups. Multi-level device groups are used to centrally manage the policies across all deployment locations with common requirements. Device group examples may be determined geographically (e.g., Europe and North America). Each device group could have a functional sub-device group (e.g., perimeter or data center). This allows for common policy sharing across different virtual systems on a device.

You can use shared policies for central control while still providing your local firewall administrator with the autonomy to make specific adjustments for local requirements. At the device group level, you can create shared policies that are defined as the first set of rules (pre-rules) and the last set of rules (post-rules) to be evaluated against match criteria. Pre- and post-rules can be viewed on a managed firewall, but they can only be edited from Panorama within the context of the administrative roles that have been defined. Local device rules (those between pre- and post-rules) can be edited by either your local firewall administrator or by a Panorama administrator who has switched to a local firewall context. In addition, an organization can use shared objects defined by a Panorama administrator, which can be referenced by locally managed device rules.

- **Role-Based Administration:** Role-based administration is used to delegate feature-level administrative access, including availability of data (enabled, read-only, or disabled and hidden from view) to different members of your staff. Specific individuals can be given appropriate access to the tasks that are pertinent to their job while making other access either hidden or read-only.

Software, Content and License Update Management: As your deployment grows in size, you may want to make sure that updates are sent to downstream boxes in an organized manner. For instance, security teams may prefer to centrally qualify a software update before it is delivered via Panorama to all production firewalls at once. Using Panorama, the update process can be centrally managed for software updates, content (application updates, antivirus signatures, threat signatures, URL filtering database, etc.), and licenses.

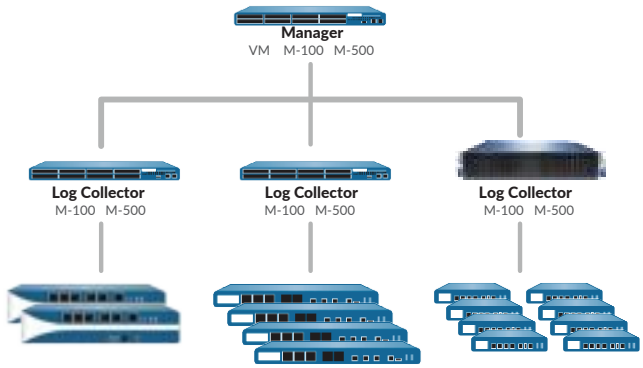


Figure 3: Panorama architecture

Using templates, device groups, role-based administration, and update management, you can delegate appropriate access to all management functions, visualization tools, policy creation, reporting and logging at a global level as well as the local level.

Deployment Flexibility

Organizations can deploy Panorama either with hardware appliances or as virtual appliances.

Hardware Appliances

Panorama can be deployed on the M-100 or the M-500 management appliances, and individual management and logging components can be separated in a distributed manner to accommodate large volumes of log data. Panorama running on these appliances can be deployed in the following ways:

- **Centralized:** In this scenario, all Panorama management and logging functions are consolidated into a single device (with the option for high availability).
- **Distributed:** You can separate the management and logging functions across multiple devices, splitting the functions between managers and log collectors.
- **Panorama Manager:** The Panorama manager is responsible for handling the tasks associated with policy and device configuration across all managed devices. The manager does not store log data locally, but rather uses separate log collectors for handling log data. The manager analyzes the data stored in the log collectors for centralized reporting.

- **Panorama Log Collector:** Organizations with high logging volume and retention requirements can deploy dedicated Panorama log collector devices that will aggregate log information from multiple managed firewalls.

The separation of management and log collection enables you to optimize your Panorama deployment in order to meet scalability, organizational or geographical requirements.

Virtual Appliance

Panorama can also be deployed as a virtual appliance on VMware® ESXi™, allowing organizations to support their virtualization initiatives and consolidate rack space, which is sometimes limited or costly in a data center.

The virtual appliance can serve as a Panorama manager and is responsible for handling the tasks associated with policy and device configuration across all managed devices. It can be deployed in two ways:

- **Centralized:** All Panorama management and logging functions are consolidated into a single virtual appliance (with the option for high availability).
- **Distributed:** Panorama distributed log collection requires a mix of the hardware and virtual appliance.

Note: The virtual appliance may not be used as a Panorama log collector. Panorama log collectors (M-100 or M-500 appliances) are responsible for offloading intensive log collection and processing tasks.

Providing the choice of either a hardware or virtualized appliance, as well as the choice to combine or separate the Panorama functions, provides you with the maximum flexibility for managing multiple Palo Alto Networks firewalls in a distributed network environment.

Panorama Specifications
Number of Devices Supported
• Up to 1,000
High Availability
• Active/Passive
Administrator Authentication
• Local database • RADIUS
Management Tools and APIs
• Graphical User Interface (GUI) • Command Line Interface (CLI) • XML-based REST API

Virtual Appliance Specifications
Minimum Server Requirements:
• 40 GB hard drive • 4 CPU cores • 4 GB RAM
VMware Support
• VMware ESX 3.5, 4.0, 4.1, 5.0
Browser Support
• IE v7 or greater • Firefox v3.6 or greater • Safari v5.0 or greater • Chrome v11.0 or greater
Log Storage
• VMware Virtual Disk: 8 TB maximum • NFS



Figure 4: M-100 Panorama Appliance



Figure 5: M-500 Panorama Appliance

M-100 Appliance
I/O
• (4) 10/100/1000, [1] DB9 console serial port, (1) USB port • Currently supported: (3) 10/100/1000, (1) DB9 console serial port
Storage
• Maximum Supported: 4 TB RAID: 8 x 1 TB RAID Certified HDD for 4 TB of RAID storage
Power Supply/Max Power Consumption
• 500W/500W
Max BTU/hr
• 1,705 BTU/hr
Input Voltage (Input frequency)
• 100-240 VAC (50-60Hz)
Max Current Consumption
• 10A @ 100 VAC
Mean Time Between Failures (MTBF)
• 14.5 years
Rack Mountable (Dimensions)
• 1U, 19” standard rack (1.75"H x 23"D x 17.2"W)
Weight
• 26.7 lbs.
Safety
• UL, CUL, CB
EMI
• FCC Class A, CE Class A, VCCI Class A
Environment
• Operating Temperature: 40° to 104° F, 5 to 40° C • Non-operating Temperature: -40° to 149° F, -40° to 65° C

M-500 Appliance
I/O
• (4) 10/100/1000, (1) DB9 console serial port, (1) USB port, (2) 10 GigE ports • Currently supported: (3) 10/100/1000, (1) DB9 console serial port
Storage
• Maximum Configuration: 12 TB RAID: 24 x 1 TB RAID Certified HDD for 12 TB of RAID storage • Currently Supported: 8 RAID pairs with 16 X 1 TB RAID Certified HDD for 8 TB of RAID storage • Default Shipping Configuration: 4 TB: 8 x 1TB RAID Certified HDD for 4 TB of RAID storage
Power Supply/Max Power Consumption
• Dual power supplies, hot swap redundant configuration • 1200W/493W (total system)
Max BTU/hr
• 1,681 BTU/hr
Input Voltage (Input frequency)
• 100-240 VAC (50-60 Hz)
Max Current Consumption
• 4.2A @ 120 VAC
Mean Time Between Failures (MTBF)
• 6 years
Rack Mountable (Dimensions)
• 2 U, 19” standard rack (3.5”H x 21”D x 17.5”W)
Weight
• 42.5 lbs.
Safety
• UL, CUL, CB
EMI
• FCC Class A, CE Class A, VCCI Class A
Environment
• Operating Temperature 50° to 95° F, 10° to 35° C • Non-operating Temperature -40° to 158° F, -40° to 65° C

4401 Great America Parkway
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

© 2016 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at http://www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

pan-ds-panorama-040116