PALO ALTO NETWORKS CERTIFIED NETWORK SECURITY ADMINISTRATOR STUDY GUIDE

July 2021

Palo Alto Networks, Inc.
www.paloaltonetworks.com

Table of Contents

Palo Alto Networks Strata Core Components

Welcome to the Palo Alto Networks PCNSA Study Guide. The purpose of this guide is to help you prepare for your PCNSA exam and achieve your PCNSA credential. This study guide is a summary of the key topic areas that you are expected to know to be successful at the PCNSA exam. It is organized based on the exam blueprint and key exam objectives.

Overview

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

Successfully passing this exam certifies that the successful candidate has the knowledge and skills necessary to implement the Palo Alto Networks Next-Generation Firewall PAN-OS® 10.0 platform in any environment. This exam does not cover other products such as Panorama, or the Prisma or Cortex suite of products.

Exam Details

- Certification Name: Palo Alto Networks Certified Network Security Administrator
- Delivered through Pearson VUE: www.pearsonvue.com/paloaltonetworks
- Exam Series: PCNSA
- Total Seat Time: 90 minutes
- Number of Items: 50
- Time for Exam Items: 80 minutes
- Format: Multiple Choice, Scenarios with Graphics, and Matching
- Language: English

Exam Domain	Weight
Palo Alto Networks Strata Core Components	17%
Device Management and Services	18%
Managing Objects	14%
Policy Evaluation and Management	26%
Securing Traffic	25%
Total	**100%**

Intended Audience

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

Qualifications

You should have two to three years' experience working in the Networking or Security industries and the equivalent of 6 months' experience working full-time with the Palo Alto Networks product portfolio.

You have at least 6 months' experience in Palo Alto Networks NGFW deployment and configuration.

Skills Required

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

Recommended Training

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

- Firewall Essentials: Configuration and Management (EDU-210)
- Digital learning - Firewall Essentials: Configuration and Management
- PCNSA Practice Test: PCNSA Practice Test

About This Document

Efforts have been made to introduce all relevant information that might be found in a PCNSA Certification Test. However, other related topics also might appear on any delivery of the exam. This document should not be considered a definitive test preparation guide but an introduction to the knowledge required, and these guidelines might change at any time without notice.

Disclaimer

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

Exam Domains and Objectives

Domain 1 – Palo Alto Networks Strata Core Components

1.1 Understand the components of the Palo Alto Networks Strata Portfolio

The Palo Alto Networks Cybersecurity Portfolio

The Palo Alto Networks product line is organized into the three principal groups outlined in the following sections.

Strata: Enterprise Security

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

Next-Generation Firewalls

Palo Alto Networks firewalls enable you to adopt best practices using application-, user-, device-, and content-based policies to minimize opportunities for attack. These next-generation firewalls are available as physical appliances, virtualized appliances, and cloud-delivered services, and all are managed consistently with Panorama. The firewalls secure your business with a prevention-focused architecture and integrated innovations that are easy to deploy and use. Palo Alto Networks Next-Generation Firewalls detect known and unknown threats, including those within encrypted traffic, using intelligence generated across many thousands of customer deployments. The firewalls reduce risks and prevent a broad range of attacks. For example, they enable users to access data and applications based on business requirements, and they stop credential theft and an attacker's ability to use stolen credentials.

With these next-generation firewalls, you can quickly create security rules that mirror business policy and are easy to maintain and adapt to your dynamic environment. They reduce response times with automated policy-based actions, and you can automate workflows via integration with administrative tools such as ticketing services or any system with a RESTful API.

The family of next-generation firewalls includes:

VM-Series: VM-Series virtual firewalls provide all the capabilities of the Palo Alto Networks next-generation hardware firewall in a virtual machine form factor so you can secure the environments that are vital for your competitiveness and innovation. Now you can leverage a single tool to safeguard cloud speed and software-defined agility by infusing segments and micro-segments with threat prevention.

The VM-Series firewalls support the following virtualization environments:

- Alibaba Cloud
- Amazon Web Services
- Cisco ACI
- Citrix NetScaler SDX
- Google CloudPlatform
- Kernel-Based Virtual Machine (KVM)
- Microsoft Hyper-V
- Microsoft Azure
- OpenStack
- Oracle Cloud Infrastructure
- VMware ESXi
- VMware NSX
- VMware vCloud Air

CN-Series: With the release of PAN-OS 10.0, Palo Alto Networks introduced the Container Native Series firewall (CN-Series) firewall. The CN-Series firewall is a containerized next-generation firewall that provides visibility and security for containerized application workloads on Kubernetes clusters. The CN-Series firewall natively integrates into Kubernetes (K8s) to provide complete Layer 7 visibility, application level segmentation, DNS security, and advanced threat protection for traffic going across trusted zones in both public cloud or data center environments.

Physical firewalls: PA-Series Next-Generation Firewalls are architected to provide consistent protection to your entire network perimeter, from your headquarters and office campus, branch offices, and data center to your mobile and remote workforce. Physical firewalls available include the PA-220, PA-800, PA-3200, PA-5200, and PA-7000 Series.

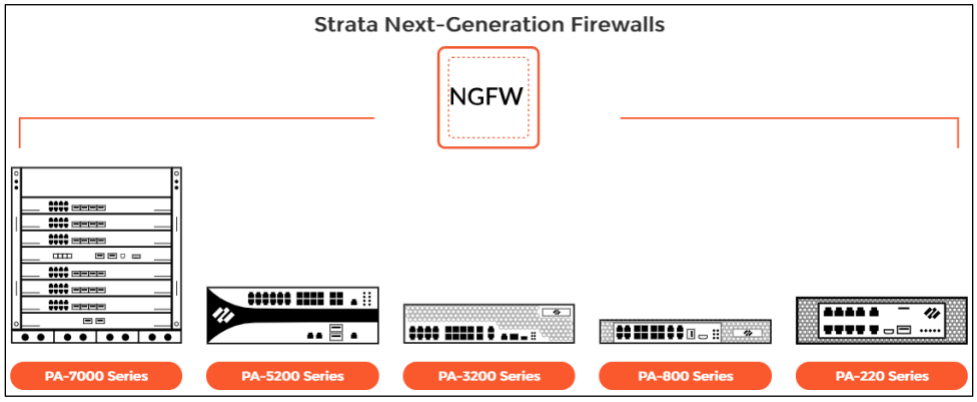

Security Subscriptions

The comprehensive range of security subscriptions extend your security policies with threat protection that is constantly kept up to date. These subscriptions are described in the following sections.

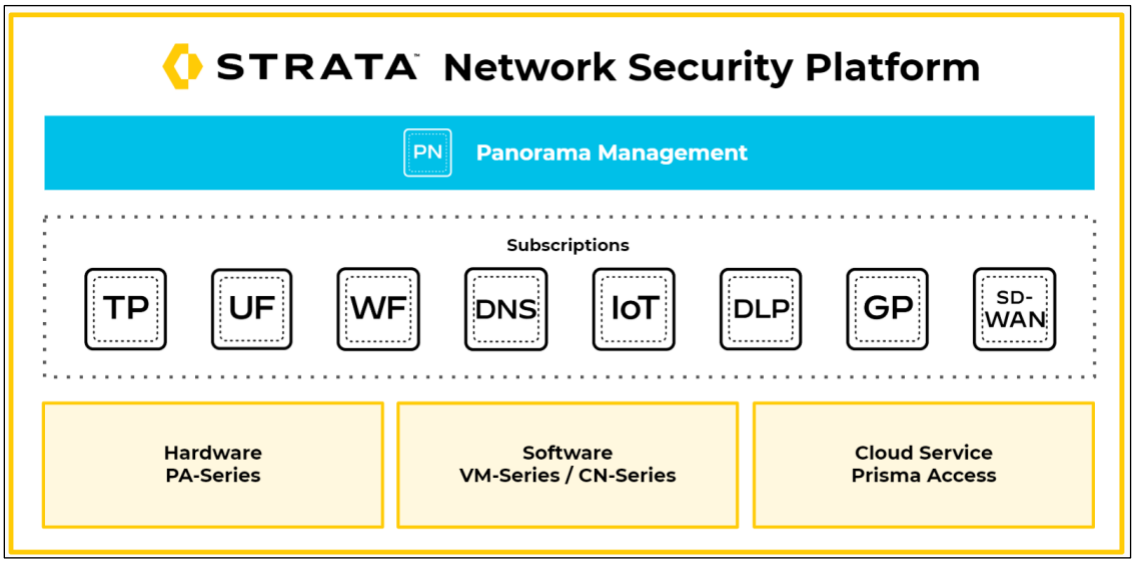

Threat Prevention: Because threats do not discriminate between application delivery vectors, an approach to security is needed that has full visibility into all application traffic, including SSL-encrypted content, with full user context. Threat Prevention leverages the visibility of the Palo Alto Networks Next-Generation Firewall to inspect all traffic and thus automatically prevents known threats regardless of port, protocol, or SSL encryption. Threat Prevention automatically stops vulnerability exploits with IPS capabilities, offers inline malware protection, and blocks outbound command-and-control traffic. When these protections are combined with WildFire® and URL filtering, owning organizations are shielded at every stage of the attack lifecycle. Protection from both known and zero-day threats is provided.

URL Filtering: Most attacks and exposure to malicious content occur during normal web browsing activities. URL filtering with PAN-DB automatically prevents attacks that leverage the web as an attack vector. These attacks include phishing links in emails, phishing sites, HTTP-based command-and-control, malicious sites, and pages that carry exploit kits. URL filtering provides the following benefits:

- Reduction of infection risk from dangerous websites and protection of users and data from malware and credential-phishing pages
- Protection across the attack lifecycle through integration with WildFire and the Cybersecurity Portfolio
- Retention of protections synchronized with the latest threat intelligence through the Palo Alto Networks cloud-based URL categorization for phishing, malware, and undesired content
- Full visibility and threat inspection into normally opaque web traffic through granular control over SSL decryption

WildFire: WildFire turns every Palo Alto Networks platform deployment into a distributed sensor and enforcement point to stop zero-day malware and exploits before they can spread and become successful. Within the WildFire environment, threats are detonated, intelligence is extracted, and preventions are automatically orchestrated across the Palo Alto Networks next-generation security product portfolio as soon as a signature is generated, thus minimizing the window in which malware can infiltrate your network. WildFire goes beyond traditional approaches. The service employs a unique, multi-technique approach that combines dynamic and static analysis, innovative machine learning techniques, and a groundbreaking bare metal analysis environment to detect unknown threats and prevent even the most evasive threats. The following illustration depicts WildFire, its information sources, and the services it supports.

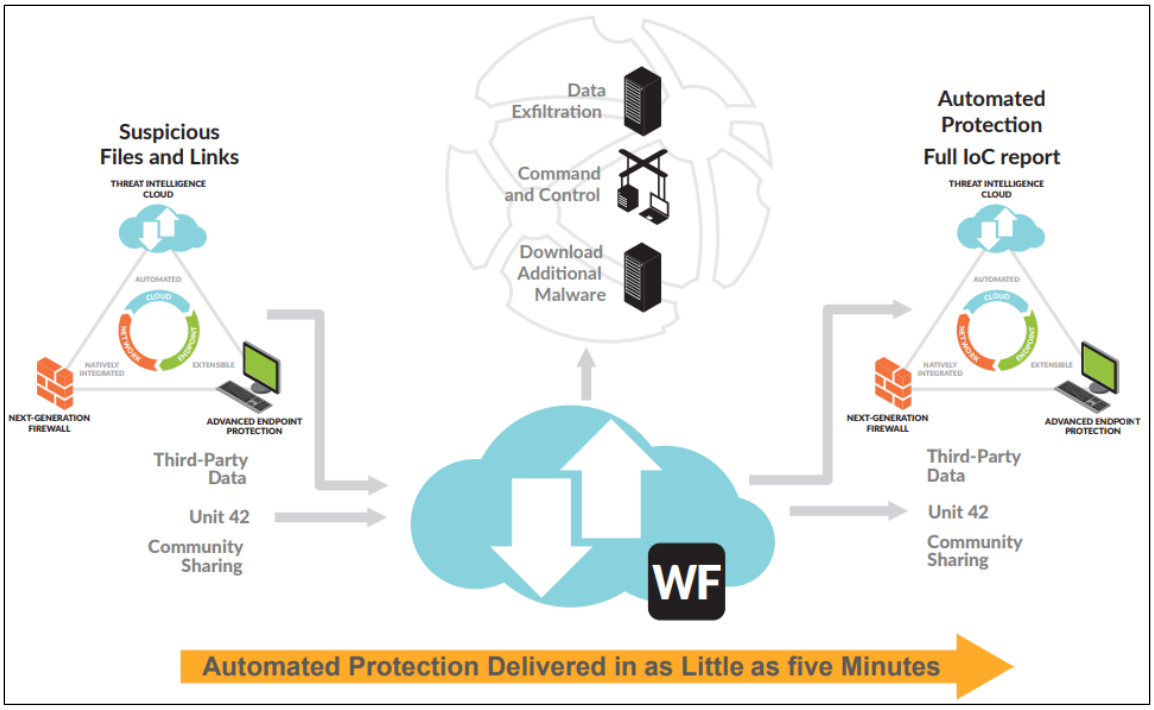

DNS Security Services: DNS Security Services applies predictive analytics, machine learning, and automation to block attacks that use DNS. Tight integration with the next-generation firewall gives you automated protections and eliminates the need for independent tools. Now you can rapidly predict and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

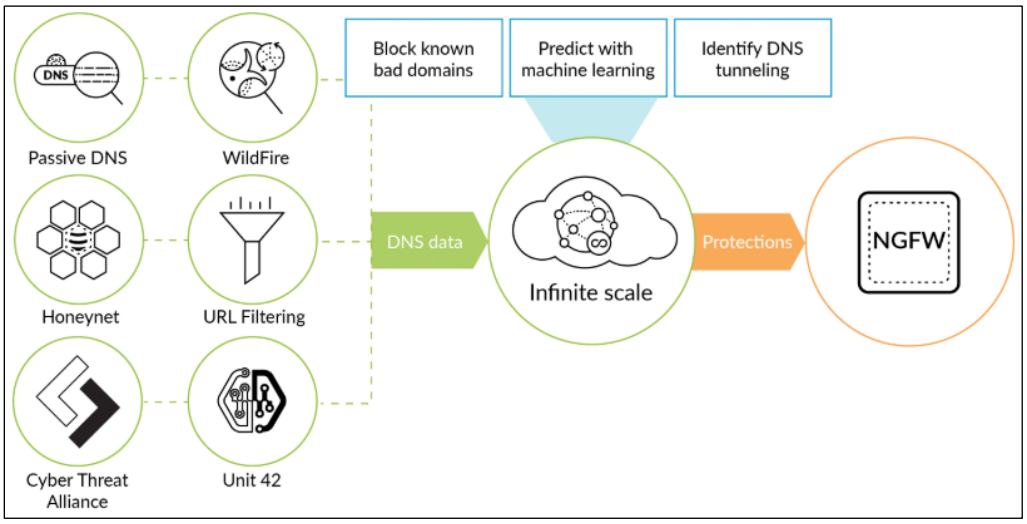

Internet of things (IoT): IoT security is a security strategy to provide complete visibility, in-depth risk analysis, and built-in enforcement. IoT security accurately identifies and classifies your IoT devices using machine learning, including devices that have never been seen before. Log files are uploaded to Cortex Data Lake and then analyzed by the IoT cloud. Log analysis enables you to understand device anomalies, vulnerabilities, and severity levels to help you make a confident decision and automate enforcement with Device-ID-enabled Security policy recommendations. The following illustration depicts this relationship.

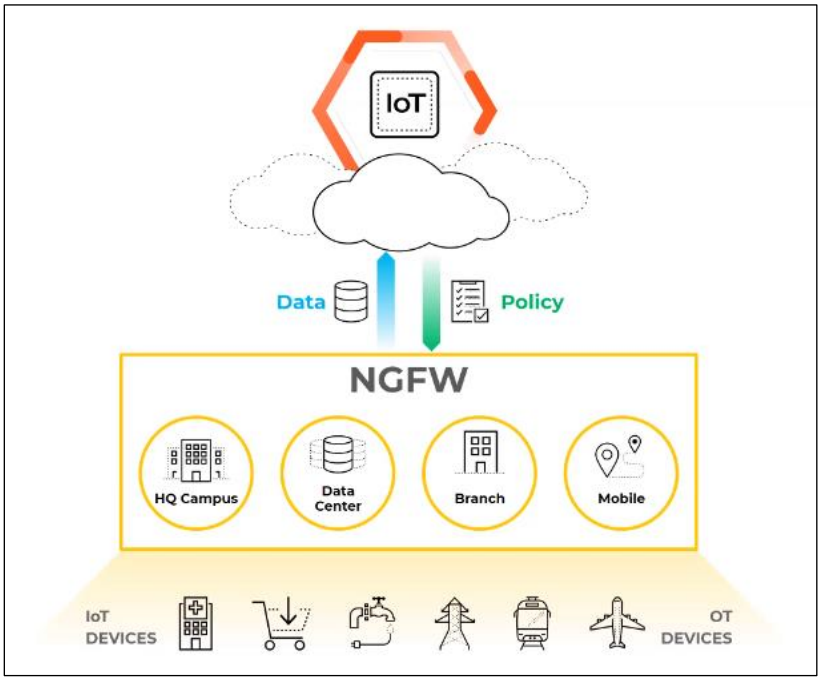

Data Loss Prevention (DLP): DLP is the practice of protecting and securing data (e.g., intellectual property, financial data, and customer or employee information) to prevent data from being lost, stolen, accessed, or misused by unauthorized individuals. DLP provides consistent data protection policy across networks, clouds, and users; accurate discovery and protection of sensitive data using automatic classification, context, and machine learning; and simplified adoption and management embedding the cloud DLP engine in existing security controls. DLP protects data in motion and data at rest via predefined data patterns and automated data profiles, supported by machine learning classifiers.

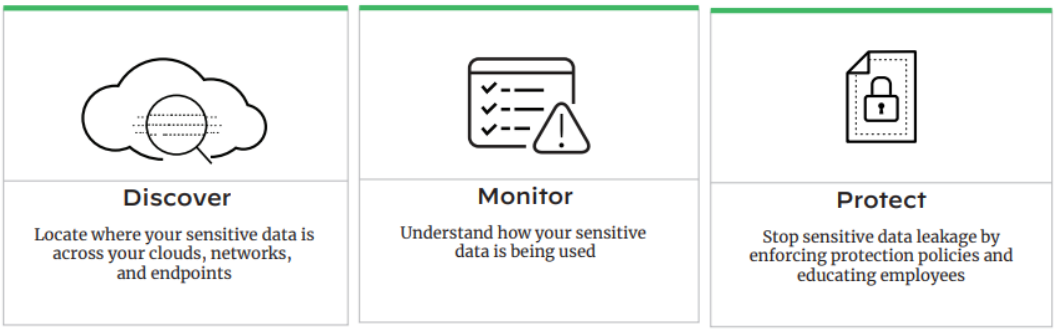

GlobalProtect: GlobalProtect safeguards your mobile workforce by inspecting all traffic using your next-generation firewalls deployed as internet gateways, whether at the perimeter, in the Demilitarized Zone (DMZ), or in the cloud. Laptops, smartphones, and tablets with the GlobalProtect app automatically establish a secure IPsec/SSL VPN connection to the next-generation firewall using the best gateway, thus providing full visibility of all network traffic, applications, ports, and protocols.

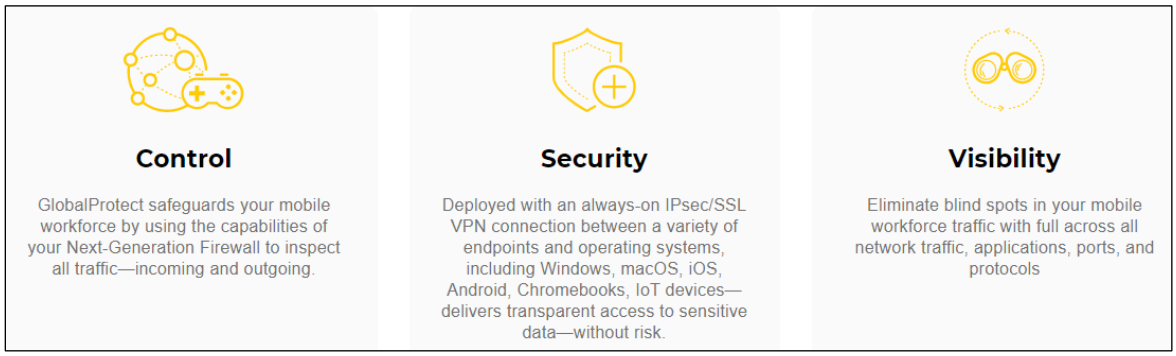

SD-WAN: This technology allows you to use multiple internet and private services to create an intelligent and dynamic WAN. The SD-WAN plugin is integrated with PAN-OS software so that you get the security features of a PAN-OS firewall and SD-WAN functionality. The SD-WAN overlay supports dynamic, intelligent path selection based on applications and services and the conditions of links that each application or service can use. Dynamic path selection avoids brownout and node failure problems because sessions can fail over to a better performing path in less than one second. The SD-WAN overlay works with all PAN-OS security features such as User-ID and App-ID to provide complete security control over your branch offices. You can configure and manage SD-WAN centrally from the Panorama web interface or the Panorama REST API.

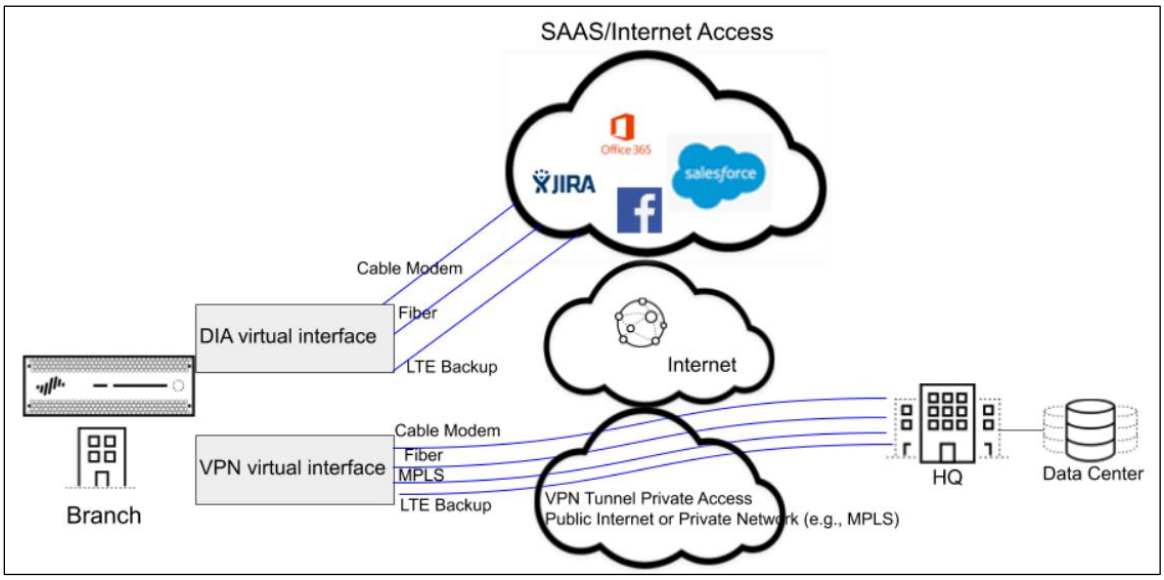

Panorama: Panorama offers easy-to-implement and centralized management features to gain insight into network-wide traffic and threats and administer your next-generation firewalls everywhere. Panorama is available in both appliance and virtual forms. Panorama provides the following features:

- Centralized management of multiple next--generation firewalls
- Software and content management of next-generation firewalls
- License management of next-generation firewalls
- Log aggregation of managed next-generation firewalls
- User-ID redistribution to managed devices
- Enterprise-level reporting
- Implementation of enterprise-level administration

The following diagram illustrates the relationship between the primary features of Panorama and its managed next-generation firewalls. The green arrow represents centralized management of firewalls. The other colored data flow arrows are labeled accordingly.

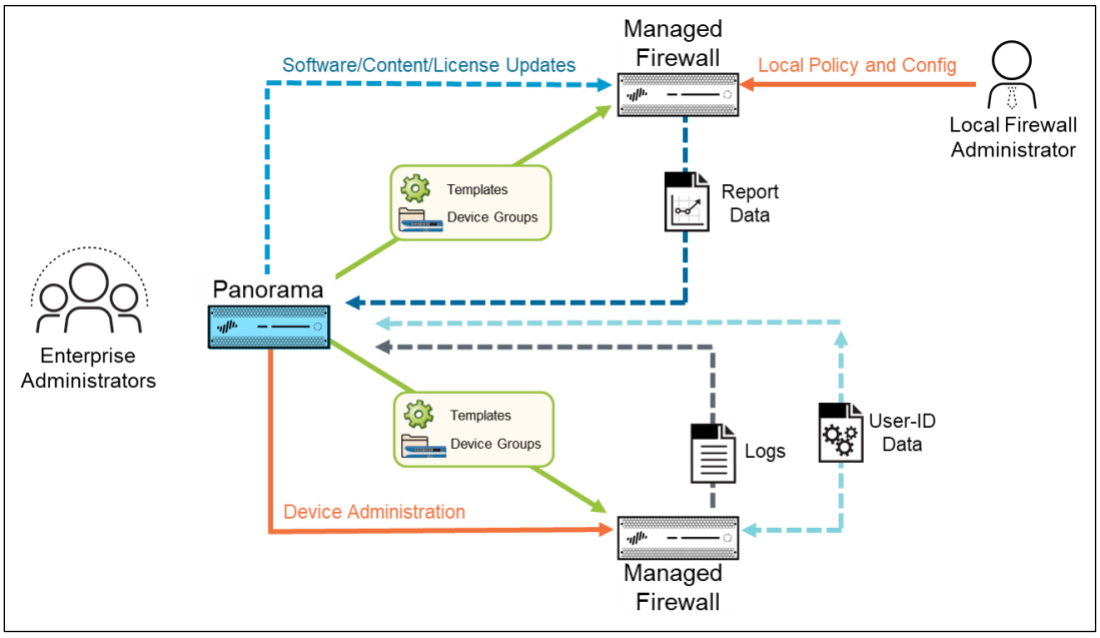

Prisma: Cloud Security

Prisma Cloud delivers complete security across the development lifecycle on any cloud, thus enabling you to develop cloud-native applications with confidence.

Prisma Cloud

Prisma Cloud is a Cloud Security Posture Management (CSPM) and Cloud Workload Protection Platform (CWPP) that provides comprehensive visibility and threat detection across your organization's hybrid, multi-cloud infrastructure.

Prisma Cloud taps into the cloud providers' APIs for read-only access to your network traffic, user activity, and configuration of systems and services, and it correlates these disparate data sets to help cloud compliance and security analytics teams prioritize risks and quickly respond to issues. It also uses an agent-based approach to secure your host, container, and serverless computing environments against vulnerabilities, malware, and compliance violations.

Cloud-native security platforms share context about infrastructure, PaaS, users, development platforms, data, and application workloads across platform components to enhance security. They also:

- Provide unified visibility for SecOps and DevOps teams
- Deliver an integrated set of capabilities to respond to threats and automate responses across any compute, network, or cloud-native application service
- Automate and enforce hundreds of out-of-the-box remediations of vulnerabilities governance policies that help ensure compliance and enforce good behavior
- Eliminate issues early and prevent alert fatigue by seamlessly integrating security early and misconfiguration alerts consistently throughout the application lifecycle, from IDE, SCM, CI/CD, and registries to runtime
- Leverage continuous vulnerability management and automated risk prioritization across the entire build-deploy-run cloud-native stack and lifecycle.
- Easily investigate any incident.
- Monitor, secure, and maintain compliance on multi- and hybrid-cloud environments with a single integrated platform
- Leverage purpose-built solutions for public clouds such as AWS, Google Cloud, and Microsoft Azure, and secure your on-premises investments such as OpenShift

Prisma Cloud secures the following cloud-native infrastructures:

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

Prisma Access (SASE)

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

SASE converges the capabilities of WAN with network security to support the needs of the digital enterprise. These disparate networks and security services include SD-WAN, secure web gateway, cloud access security broker (CASB), software-defined perimeter, DNS protection, and firewall as a service.

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

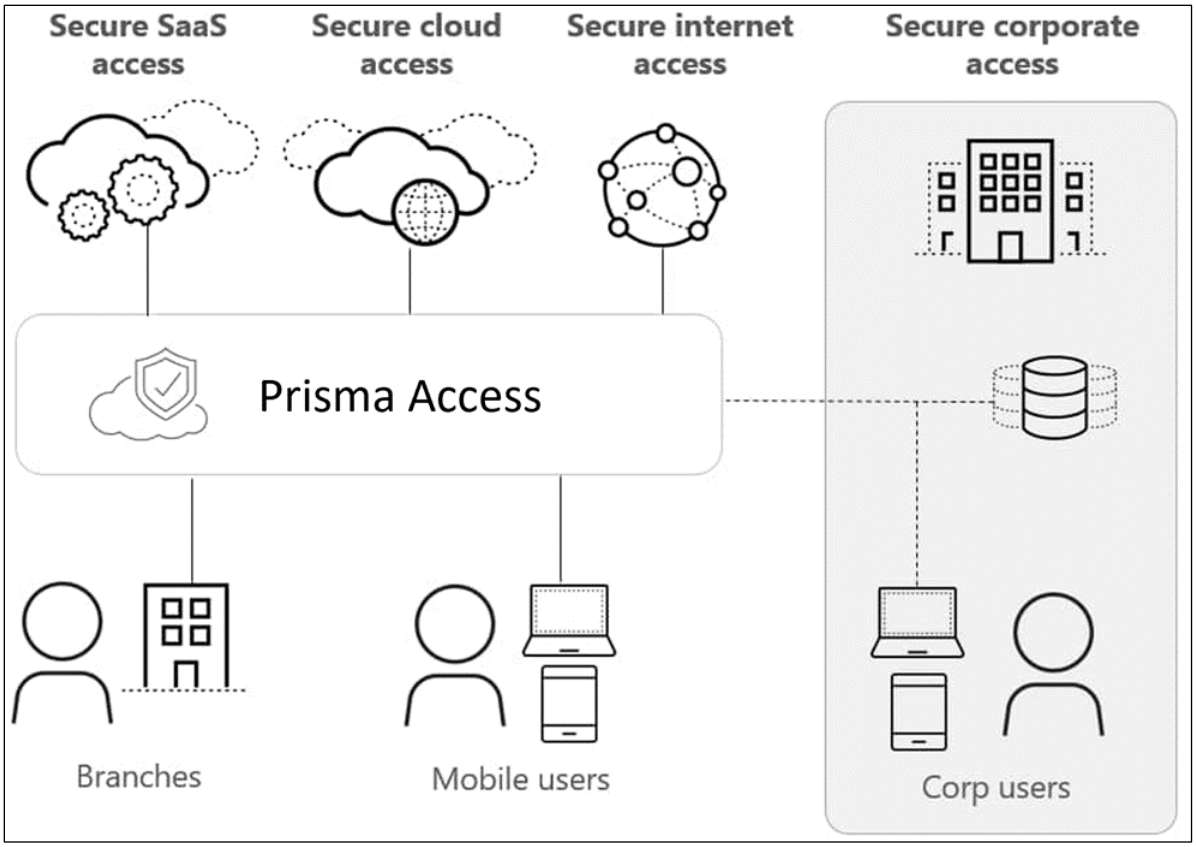

Prisma SaaS

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

VM-Series Next-Generation Firewalls

VM-Series is the virtualized form factor of the Palo Alto Networks Next-Generation Firewall. To meet the growing need for inline security across diverse cloud and virtualization use cases, you can deploy the VM-Series firewall on a wide range of private and public cloud computing environments. See the description for the “Strata: Enterprise Security” in the “Next-Generation Firewalls” section.

Cortex: Security Operations

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

Cortex XDR

The Cortex XDR app offers you complete visibility over network traffic, user behavior, and endpoint activity. It simplifies threat investigation by correlating logs from your ***sensors*** to reveal threat causalities and timelines. These summarizations enable you to easily identify the root cause of every alert. The app also allows you to perform immediate response actions. Finally, to stop future attacks, you can proactively define **indicators of compromise (IoCs)** and BIOCs to detect and respond to malicious activity. The following illustration depicts the Cortex XDR architecture.

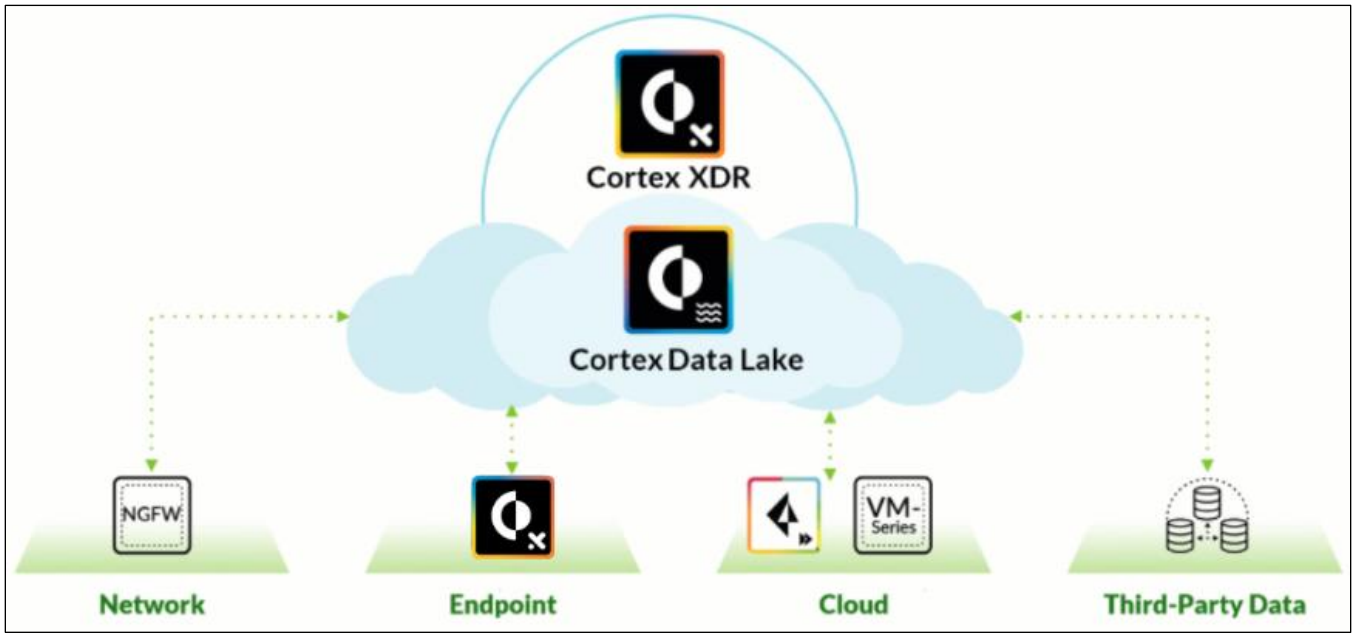

Cortex XSOAR

Cortex XSOAR is the industry-leading Security Orchestration, Automation, and Response (SOAR) technology that will automate up to 95% of all response actions requiring human review and allow overloaded security teams to focus on the actions that really require their attention. Cortex SOAR integrates with a wide variety of products providing enhanced automation and response across processes involving multiple products. The following illustration depicts the Cortex XSOAR engine in the center with information sources on the left and potential consumers on the right.

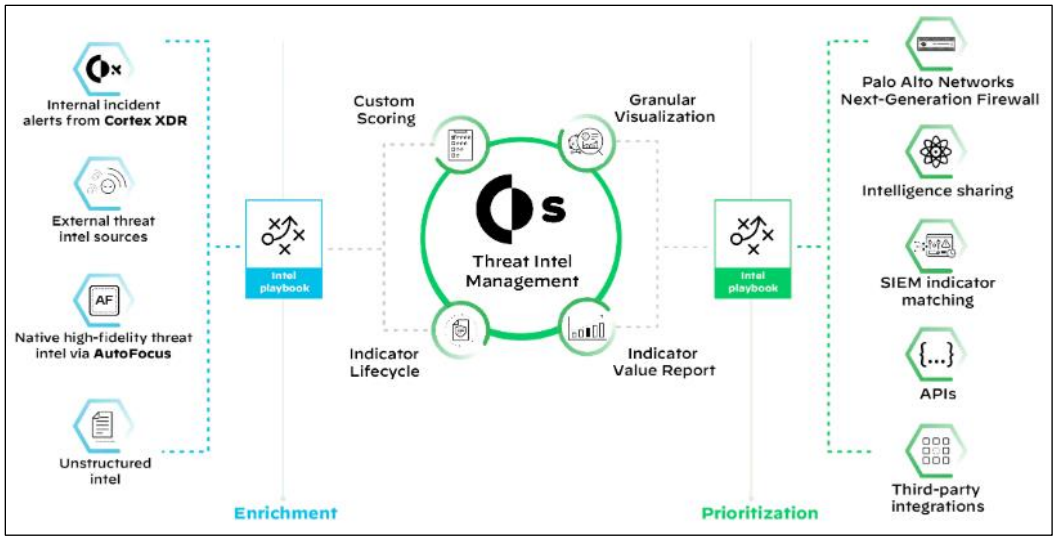

Cortex Data Lake

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

- Radically simplify your security operations by collecting, integrating, and normalizing your enterprise's security data
- Effortlessly run advanced AI and machine learning with cloud-scale data and compute
- Constantly learn from new data sources to evolve your defenses

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

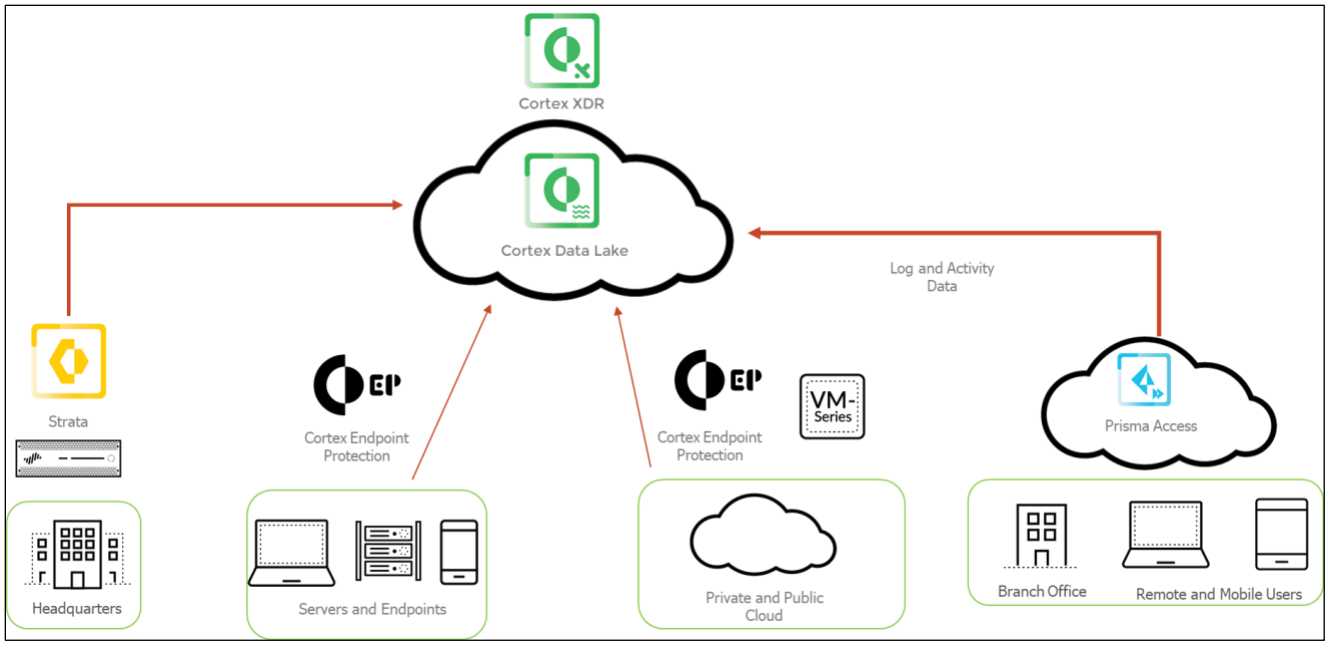

Following are the products that use Cortex Data Lake and their requirements:

- Palo Alto Networks Next-Generation Firewalls and Prisma Access:
    - Next-generation firewalls and Panorama for network security management with the ability to connect to the cloud service
    - Next-generation firewalls and Panorama running PAN-OS® 8.0.5+
    - Panorama with the cloud services plugin installed
- Old versions of Palo Alto Networks Traps for endpoint protection and response:
    - Traps running version 5.0+ with Traps management service
- Cortex XDR:
    - Cortex XDR application (Traps agent included)

AutoFocus

AutoFocus contextual threat intelligence service is your single source for threat intelligence. Your teams will receive instant understanding of every event with context from Unit 42 threat researchers, and you can embed rich threat intelligence in analysts' existing tools to significantly speed investigation, prevention, and response.

- Get unique visibility into attacks crowdsourced from the industry's largest footprint of network, endpoint, and cloud intel sources.
- Enrich every threat with the deepest context from Unit 42 threat researchers.
- Give analysts a major time advantage with intel embedded in any tool through a custom threat feed and agile APIs.

The following illustration depicts AutoFocus as the repository for many information sources within Palo Alto Networks and externally.

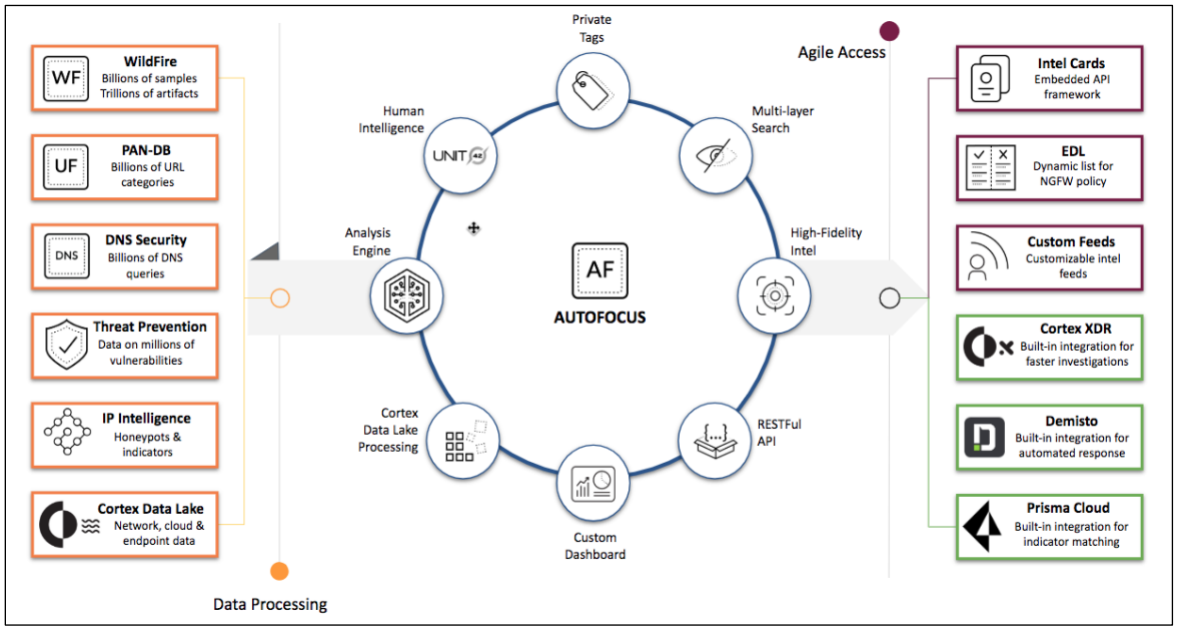

Sample Questions

Q1. What are four components of the Palo Alto Networks Cybersecurity Portfolio? (Choose four.)

a) Cortex DynamicDNS

b) WildFire

c) Cortex XDR

d) OpenConnect

e) Prisma Access

f) AutoFocus

Q2. Which cloud-delivered security service provides instant access to community-based threat data?

a) Prisma SaaS

b) AutoFocus

c) Unit 42

d) Cortex XDR

Q3. Which cloud-delivered security service provides security and connectivity for branches and mobile users?

a) Cortex XSOAR

b) Cortex XDR

c) AutoFocus

d) Prisma Access

Q4. Which Palo Alto Networks cybersecurity portfolio product provides access to applications from Palo Alto Networks, third parties, and customers?

a) WildFire

b) Cortex Data Lake

c) Network Security

d) Prisma Access

Q5. Which Palo Alto Networks firewall feature provides all the following abilities?

- Stops malware, exploits, and ransomware before they can compromise endpoints
- Provides protection while endpoints are online and offline, on network and off
- Coordinates enforcement with network and cloud security to prevent successful attacks
- Detects threats and automates containment to minimize impact
- Creates zero-day malware signatures with cloud-based threat analysis
- Integrates with Palo Alto Networks Cortex Data Lake

a) Cortex XDR

b) Prisma SaaS

c) WildFire

d) AutoFocus

Identify stages in the cyberattack lifecycle

The cyberattack lifecycle is a sequence of events that an attacker goes through to successfully infiltrate a network and exfiltrate data from it. Blockage of just a single stage in this lifecycle often is enough to protect a company's network from a successful attack. Palo Alto Networks products prevent advanced cyberattacks at every stage of the attack lifecycle. The Palo Alto Networks platform protects every part of the global enterprise network: It addresses vulnerabilities and malware arriving at the endpoint, mobile device, and network perimeter, or within the data center.

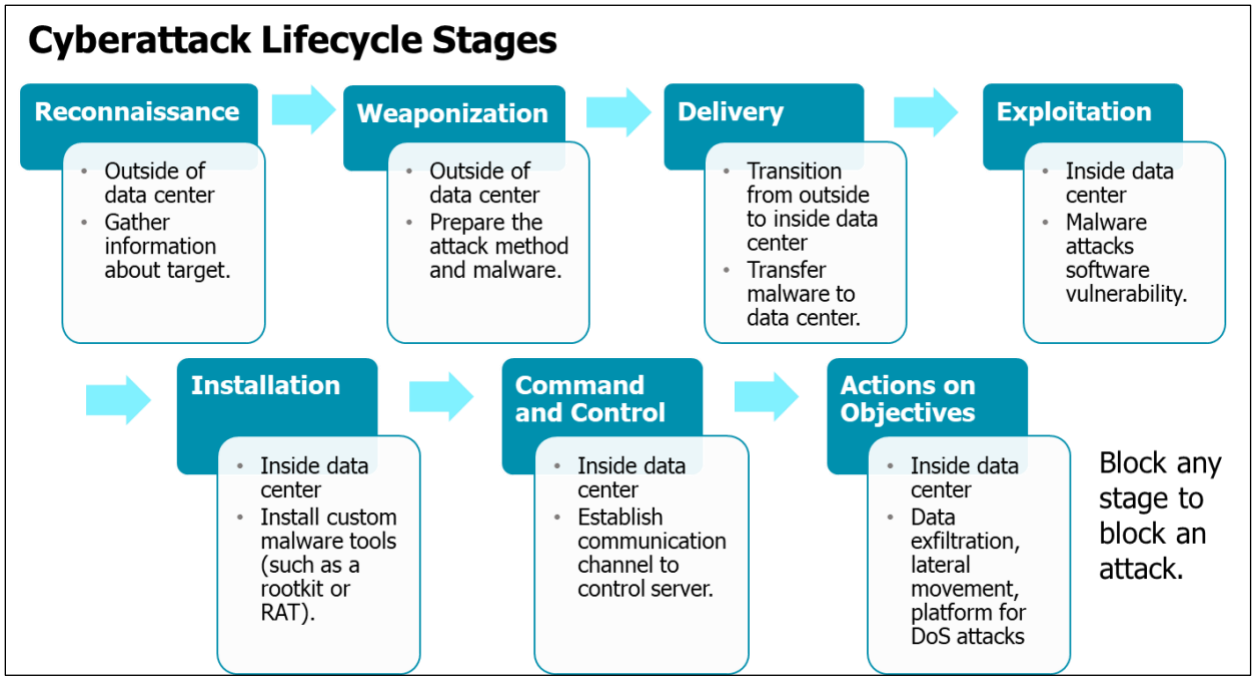

When cyberattackers strategize their way to infiltrate an organization's network and exfiltrate data, they follow the series of stages that comprise the attack lifecycle. They must progress through each stage to successfully complete an attack. Block cyberattacks at any point in the cycle to break the chain of attack. Note that the attacks can follow any order with the attack chain. The following sections describe the different stages of the attack lifecycle and steps that should be taken at each stage to prevent an attack.

1. **Reconnaissance:** During the first stage of the attack lifecycle, cyber adversaries carefully plan their method of attack. Attackers research, identify, and select targets within an organization such as human resources and financial personnel that will allow them to meet their objectives. Attackers can gather intelligence through publicly available sources such as Twitter, LinkedIn, and corporate websites, all the places where a company will share information about itself. Cyberattackers also will scan for vulnerabilities that can be exploited within the target network (services and applications) and map out areas that they can take advantage of.

Prevent by:

- Performing continuous inspection of network traffic flows to detect and prevent port scans and host sweeps
- Implementing security awareness by limiting what should be posted on the internet. Examples of content that should not be posted are sensitive documents, customer lists, event attendees, job roles, and responsibilities

2. **Weaponization:** If any vulnerability has been detected by reconnaissance, attackers next determine which methods to use to deliver malicious payloads. Methods they might use include automated tools such as exploit kits, spear phishing attacks with malicious links, infected attachments, and malvertizing. All Weaponization activity occurs on machines away from the target. No prevention steps are possible or required at this point.
3. **Delivery:** This stage marks the transition from the attacker working outside of an organization's network to working within an organization's network. Malware delivered during this stage is designed to exploit existing software vulnerabilities. To deliver its initial malware, the attacker might choose to embed malicious code within seemingly innocuous PDF or Word files, or within an email message. For highly targeted attacks, an attacker might craft a deliverable related to the specific interests of an individual that might entice the individual into accessing a malicious website or opening an infected email message.

Prevent by:

- Gaining full visibility into all traffic, including SSL traffic, by decrypting it and blocking high-risk applications
- Extending protections to remote and mobile devices
- Protecting against perimeter breaches by blocking malicious or risky websites using URL filtering
- Blocking known exploits, malware, and inbound C2 communications using multiple threat prevention disciplines, including IPS, anti-malware, anti-C2, DNS monitoring, sinkholing, and file and content blocking
- Detecting unknown malware and automatically informing customers and third parties globally to thwart new attacks
- Providing ongoing education to users about spear phishing links, watering hole attacks, unknown emails, risky websites, malicious USB drives, and other attack methods

4. **Exploitation:** In this stage, attackers deploy an exploit against a vulnerable application or system, typically using an exploit kit or weaponized document such as a Microsoft Word .doc or Adobe

Acrobat .pdf file. An exploit kit or weaponized document enables the attacker to gain an initial entry point into the organization.

Prevent by:

- Keeping systems patched
- Educating users to recognize phishing attempts
- Blocking known and unknown vulnerability exploits on the endpoint
- Automatically delivering new protections globally to thwart follow-up attacks

5. **Installation**: After cyberattackers have established an initial foothold, they will install malware to conduct further operations, such as maintaining access, maintaining persistence, and escalating privileges. Off-the-shelf tools are the most common method of attack.

Prevent by:

- Preventing malware installation on the endpoint, network, and cloud services
- Establishing secure security zones with strictly enforced user access controls that provide ongoing monitoring and inspection of all traffic between zones (Zero Trust model)
- Limiting local administrator access for users
- Training users to identify the signs of a malware infection and know how to follow up if something occurs

6. **Command and Control:** With malware installed, attackers own both sides of the connection: their malicious infrastructure and the infected endpoint. They can actively control the system and proceed to the next stages of an attack. Attackers will establish a command channel to be able to communicate and pass data back and forth between the infected devices and their own infrastructure. Typical surveillance methods include key logging, audio capture, screen capture, and webcam capture.

Prevent by:

- Blocking outbound C2 communications
- Blocking uploads that match file and data pattern uploads
- Redirecting malicious outbound communication to internal sinkholes to identify and block compromised hosts
- Blocking outbound communication to known malicious URLs through URL filtering

- Creating a database of malicious domains to ensure global awareness and prevention through DNS monitoring
- Limiting the attacker's ability to move laterally within a network

7. **Actions on the Objective:** These actions are completed by an active attacker. After attackers have control, persistence, and ongoing communication between the endpoint and the attacker's infrastructure, they will act to achieve their goal. Their objective could be to exfiltrate data, destroy critical infrastructure, deface a website, or create fear or the means for extortion.

Prevent by:

- Using threat intelligence tools to proactively hunt for indicators of compromise (IoCs) on the network
- Monitoring and inspecting all traffic between security zones
- Enforcing user access controls across secure zones
- Blocking outbound C2 communications along with traffic that matches file and data pattern uploads
- Using URL filtering to block outbound communication to known malicious URLs
- Implementing granular control of applications and applying user control to enforce file transfer application policies on the enterprise
- Eliminating known archiving and transfer tactics and limiting the attacker's ability to move laterally within a network

As was mentioned, advanced attacks are very complex because an adversary can succeed only by progressing through every stage of the attack lifecycle. If they cannot successfully take advantage of vulnerabilities, then they cannot install malware and will not be able to obtain command and control over the system. Cybersecurity is asymmetric warfare: An attacker must do everything correctly to succeed, but a network defender needs to do only one thing correctly among multiple opportunities to prevent an attack.

Disruption of the attack lifecycle relies not only on technology but also on people and processes in the organization. The people must receive ongoing security awareness training and be educated in best practices to minimize the likelihood of an attack progressing past the first stage. Processes and policies must be in place for remediation if an attacker successfully progress through the entire attack lifecycle.

Here is a link to an on-demand webinar that examines the anatomy of real attacks carried out by advanced adversaries:

https://www.paloaltonetworks.com/resources/webcasts/defeat-pragmatic-adversary.html

Sample Questions

Q1. True or false: Blockage of just one stage in the cyberattack lifecycle will protect a company's network from attack.

a) True

b) false

Q2. What are two stages of the cyberattack lifecycle? (Choose two.)

a) weaponization and delivery

b) manipulation

c) extraction

d) command and control

Q3. Command and control can be prevented through which two methods? (Choose two.)

a) exploitation

b) DNS Sinkholing

c) URL filtering

d) reconnaissance

Q4. Exploitation can be mitigated by which two actions? (Choose two.)

a) keeping systems patched

b) using local accounts

c) blocking known and unknown vulnerability exploits on the endpoint

d) providing admin credentials

Identify the correct order of the policy evaluation based on the packet flow architecture

Policies

Palo Alto Networks firewalls implement several types of policies:

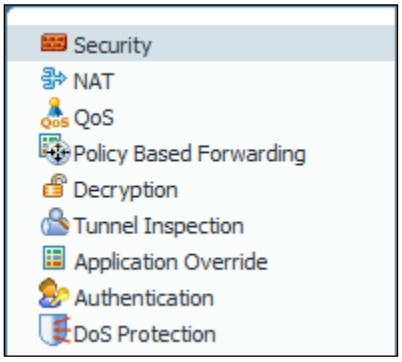

Types of Policies in a Palo Alto Networks Firewall

Every policy type contains a list of rules to match. For every connection, policy rules are matched from the top down and the first rule matched is applied, and that is the only rule from the policy that is applied to that connection. The order in which policy types are applied is based on the packet processing order:

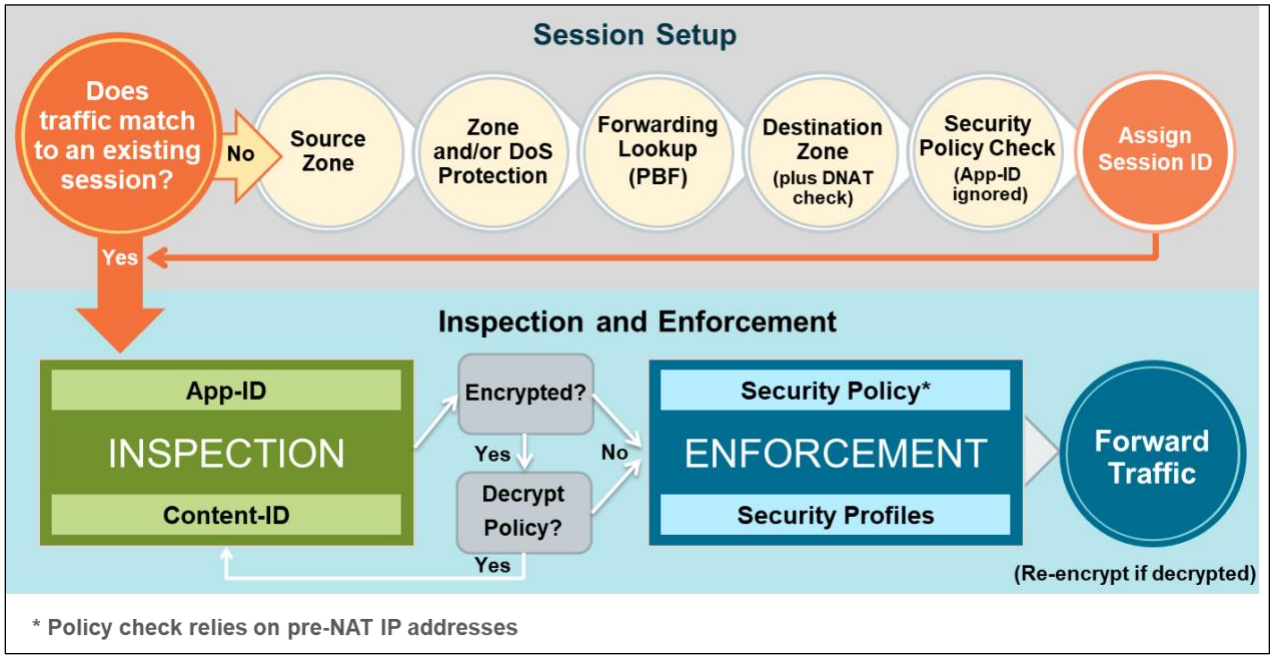

All traffic processed by the firewall follows this sequence of events.

Evaluation Order

An example of the importance of evaluation order can be found with NAT and Security policies. NAT policy rules change IP addresses in packet headers. Security policy rules are required to allow the traffic in question to transit the firewall. The processing order indicates that addresses changed by NAT policy rules are done *after* Security policy rules are evaluated, resulting in Security policy rules being re-evaluated for pre-NAT packet addresses.

An overview of the different policy types is here:
https://docs.paloaltonetworks.com/pan-os/10-0/pan-os-admin/policy/policy-types.html

Policy Match and Connectivity Tests

In PAN-OS, you can perform policy match and connectivity tests for firewalls from the web interface rather than the CLI. You can easily test the running configuration of your firewalls and verify traffic and connectivity to ensure that policy rules are matching traffic as expected to allow or deny traffic, and that firewalls can connect to network resources and external services such as WildFire®, Log Collectors, or the Content Distribution Network.

Details about using the management web interface for testing can be found here:
https://docs.paloaltonetworks.com/pan-os/10-0/pan-os-admin/policy/test-policy-rule-traffic-matches.html

Sample Questions

Q1. What is the correct order of evaluation between the Security policy and the NAT policy?

a) NAT policy evaluated, Security policy evaluated, NAT policy applied, Security policy applied

b) NAT policy evaluated, NAT policy applied, Security policy evaluated, Security policy applied

c) NAT policy evaluated, Security policy evaluated, Security policy applied, NAT policy applied

d) Security policy evaluated, NAT evaluated, NAT policy applied, Security policy applied

Q2. Which two statements are true regarding firewall policy? (Choose two.)

a) All policy rules are evaluated, and the most specific rule will match.

b) Policy rules are evaluated from the top down, and the first rule matched processes the traffic.

c) Interzone traffic is allowed by default.

d) Intrazone traffic is allowed by default.

e) Outbound traffic is allowed by default. Only inbound traffic is evaluated.

Q3. Which firewall operation order is correct?

a) decryption, check allowed ports, App-ID identification, check Security policy

b) decryption, App-ID identification, check allowed ports, check Security policy

c) check allowed ports, decryption, App-ID identification, check Security policy

d) decryption, App-ID identification, check Security policy, check allowed ports

1.2 Identify the components and operation of Single-Pass Parallel Processing architecture

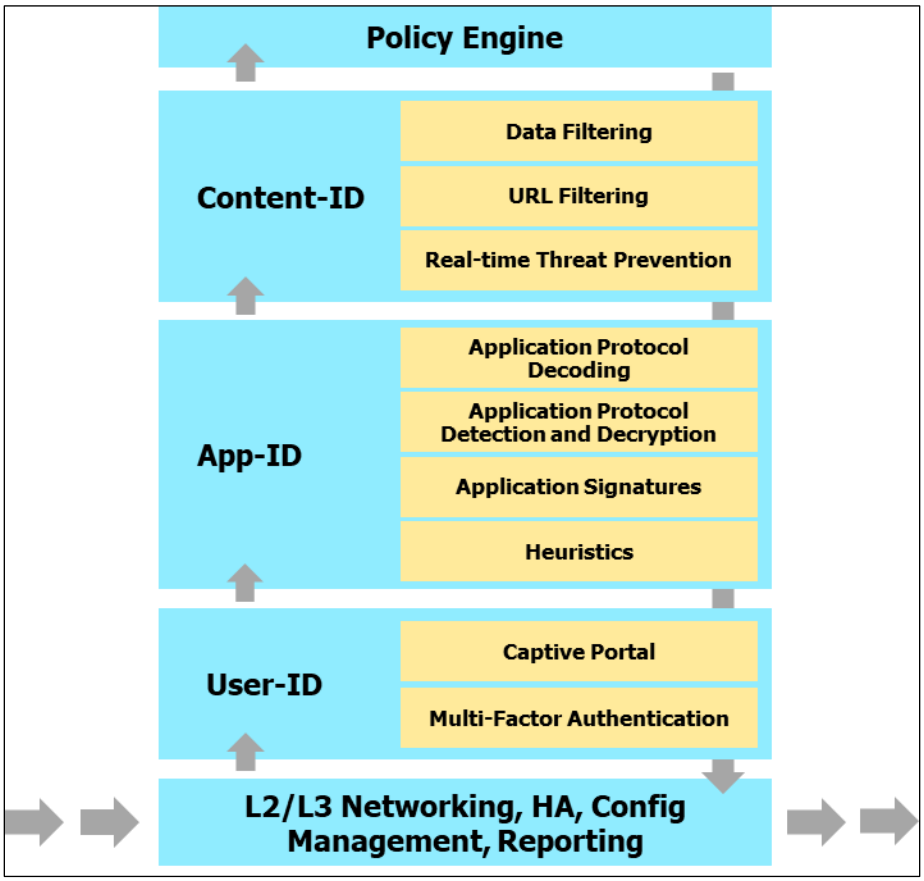

Palo Alto Networks has reduced latency enormously using the Single-Pass Parallel Processing (SP3) architecture, which combines two complementary components:

- Single-pass software
- Parallel processing hardware

The SP3 architecture is the overall design approach for Palo Alto Networks Next-Generation Firewalls. The architecture enables full, contextual classification of traffic, followed by a set of enforcement and threat prevention options. The architecture classifies and controls traffic in a “single pass” through the firewall using a variety of stream-based technology components. Each current protection feature in the

device (antivirus, spyware, data filtering, and vulnerability protection) uses this stream-based signature format. The stream-based design of the architecture results in superior performance, especially when multiple security functions are enabled.

This architecture enables you to achieve superior security posture and efficiency. The SP3 architecture enables Palo Alto Networks Next-Generation Firewalls to exceed competitors' firewall performance. Frequent design approaches used by competitors often add next-generation features in a sequence of separate engines that limit policy flexibility, negatively impact performance, and increase management complexity.

The software's "scan it all, scan it once" approach enables superior security posture and performance on both physical and virtual next-generation firewalls. The architecture incorporates advanced technologies (e.g., App-ID, User-ID, and WildFire) to provide superior classification and control capabilities to help secure your organization's network.

Management and Data Planes

In addition to the single-pass software, hardware is the other critical piece of the Palo Alto Networks SP3 architecture. The management plane and data-plane functionality on both physical and virtual firewalls is integral to all Palo Alto Networks firewalls. These separate planes have dedicated hardware resources (CPU, RAM, and storage), which makes them independent of each other. This separation means that heavy use of one plane will not adversely impact the other plane's performance. For example, an administrator could be running a very processor-intensive report, and yet the ability to process packets would not be affected by this reporting job because of the separation of the data and control planes.

The control plane provides the following management features of the firewall:

- Firewall configuration
- Logging
- Reporting

The data plane provides the following data processing features of the firewall:

- Signature matching
- Security processing
- Network processing

The following illustration depicts the architecture of a Palo Alto Networks Next-Generation Firewall.

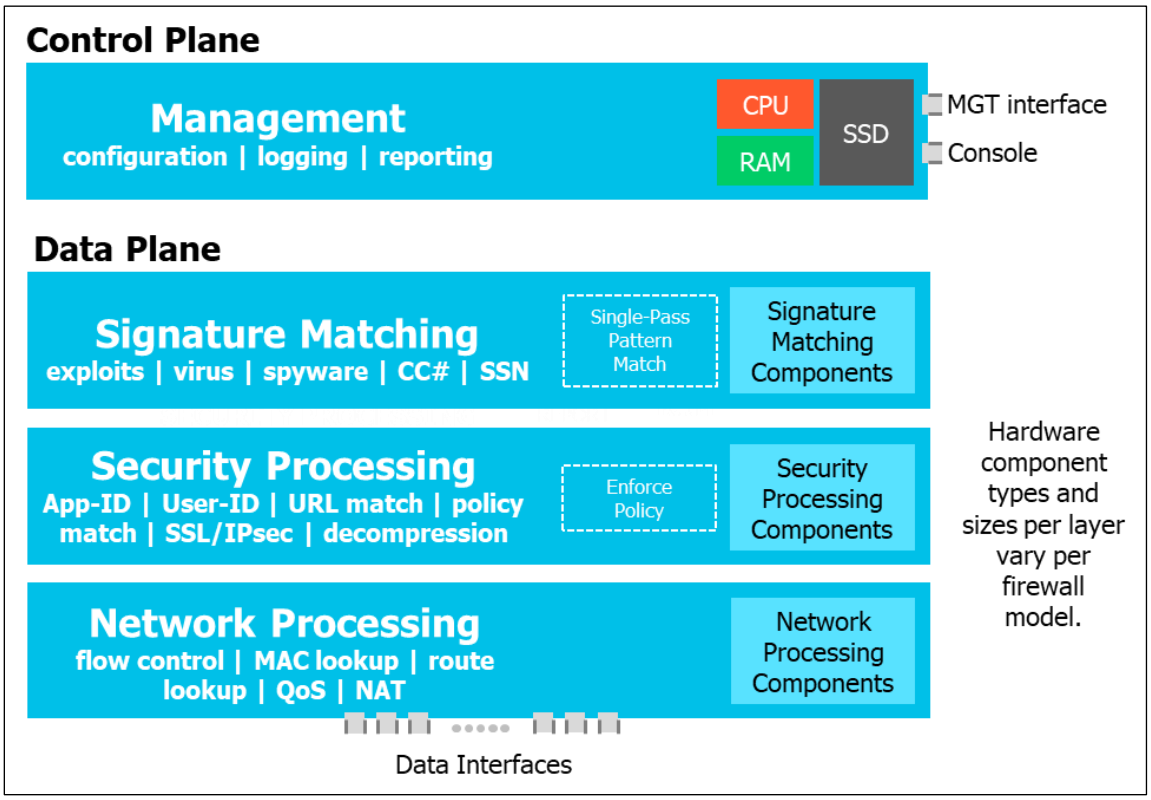

Sample Questions

Q1. Which three management features does the control plane provide? (Choose three.)

a) security processing
b) logging
c) reporting
d) firewall configuration
e) signature matching
f) network processing

Q2. Which three data processing features does the data plane provide? (Choose three.)

a) security processing
b) logging
c) reporting
d) firewall configuration
e) signature matching
f) network processing

Q3. What are three components of the Network Processing module? (Choose three.)

a) QoS
b) NAT
c) App-ID
d) flow control
e) url match
f) spyware

Q4. Which approach most accurately defines the Palo Alto Networks SP3 architecture?

a) prioritize first
b) sequential processing
c) scan it all, scan it once
d) Zero Trust segmentation platform

Q5. What is the result of using a stream-based architectural design?

a) superior performance
b) increased latency
c) detailed logging
d) increased functionality

Domain 2 – Device Management and Services

2.1 Identify and use firewall management interfaces

Management Access to the Palo Alto Networks Firewalls

Four methods are used to manage the Palo Alto Networks Next-Generation Firewalls:

- Web interface
- CLI
- Panorama
- XML API

All Palo Alto Networks firewalls provide an out-of-band management (MGT) port that you can use to perform firewall administration functions. The MGT port uses the control plane, thus separating the management functions of the firewall from the network traffic processing functions (data plane). This separation between the control plane and data plane safeguards access to the firewall and enhances performance. When you use 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. A serial/console port also is available to accomplish initial configuration of the firewall using SSH or Telnet.

Some management tasks, such as retrieving licenses and updating the threat and application signatures on the firewall, require access to the internet, which typically is done via the MGT port. If you do not want to enable external access via your MGT port, you can set up an in-band data port on the data plane to provide access to required external services (using service routes). Service routes are explained in more detail in a following section.

Initial Steps to Gain Access to the Firewall

The initial step to gaining access to the firewall for the first time is to gather the following information for the MGT port. Note that if the firewall is set up as a DHCP client, this information will be included automatically via DHCP:

- IP address
- Netmask
- Default gateway
- At least one DNS server address

The second step is to connect a computer to the firewall using either an RJ-45 Ethernet cable or a serial cable.

An RJ-45 Ethernet cable is connected from your computer to the firewall MGT port. From a browser, navigate to https://192.168.1.1. Note that you might need to change the IP address on your computer to an address in the 192.168.1.0/24 subnet, such as 192.168.1.2, to access this URL.

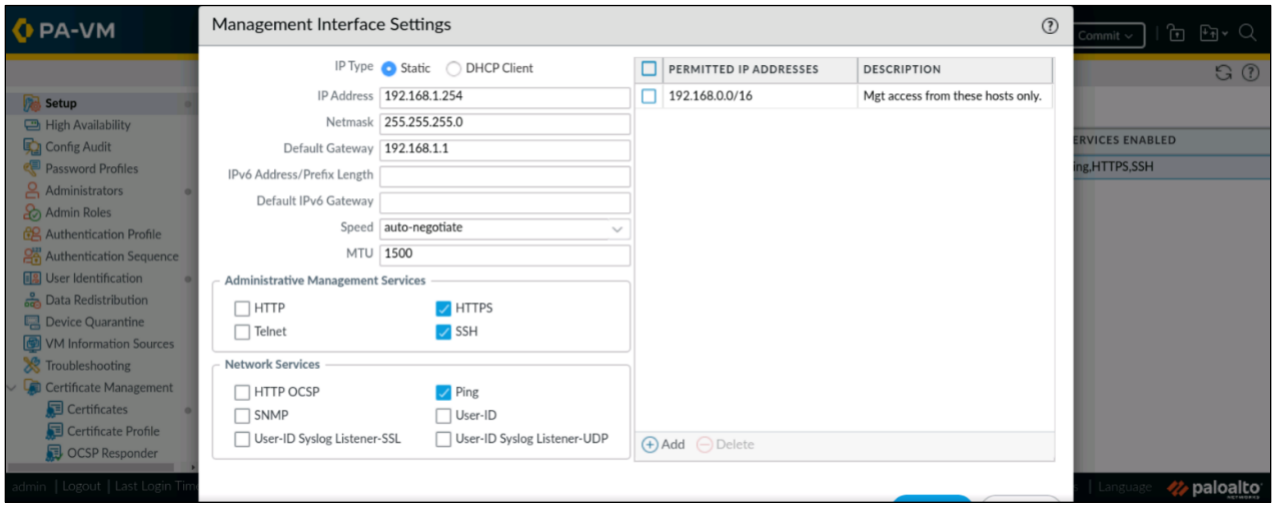

If you want to perform your initial configuration via the CLI or don't know the address served to the management port via DHCP to access the web interface, connect the serial cable from your computer to the firewall console port using terminal emulation software such as SSH or Telnet. The default connection parameters are 9600-8-N-1.

The third step is to log in to the firewall. The default username is "admin" and the default password is "admin". Starting with PAN-OS 9.1, you are forced to change the admin account password the first time you log in to the web interface.

Four Firewall Management Methods

Web interface: The web interface is used for configuration and monitoring over HTTP or HTTPS using a web browser. HTTPS is the default method; HTTP is available as a less secure method than HTTPS.

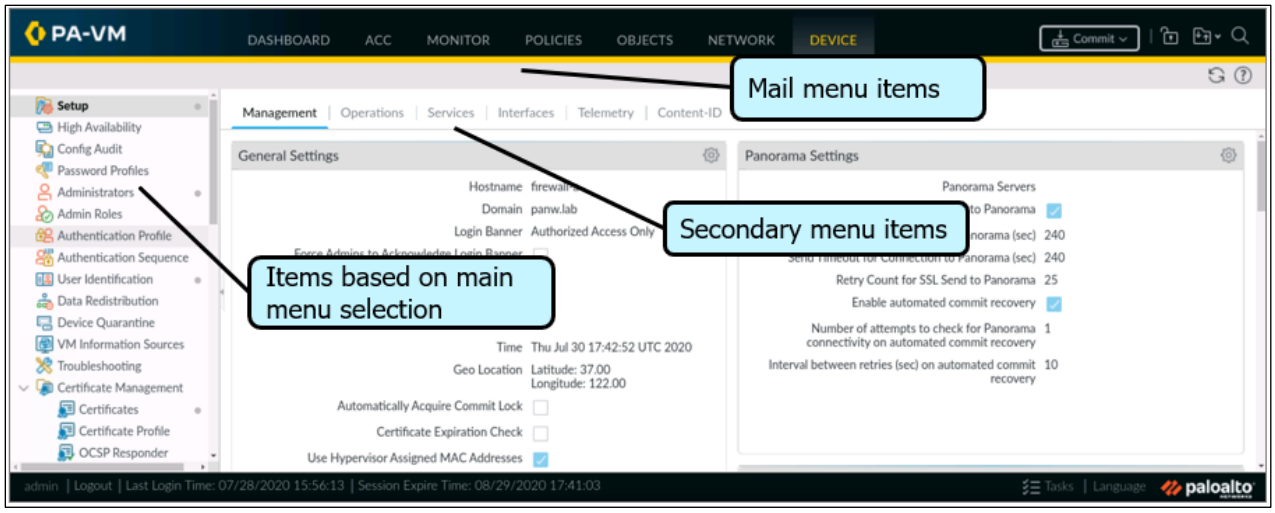

CLI: The CLI is text-based configuration and monitoring over the serial console port, or over the MGT port using SSH or Telnet. The Palo Alto Networks firewall CLI offers access to debugging information and often is used by experienced administrators for troubleshooting. The account used for authenticating into the CLI must have CLI access enabled.

The CLI command prompt will be in *operational mode* by default. The commands available within the context of operational mode include basic networking commands such as **`ping`** and **`traceroute`**, basic system commands such as **`show`**, and more advanced system commands such as **`debug`**. Commands to shut down and restart the system also are available from within operational mode.

Access configuration mode by typing the command **`configure`** while in operational mode.

Configuration mode enables you to display and modify the configuration parameters of the firewall, verify candidate configuration, and commit the config.

The following figure shows an example CLI screen with the first lines of **`show system state`** while in operational mode:

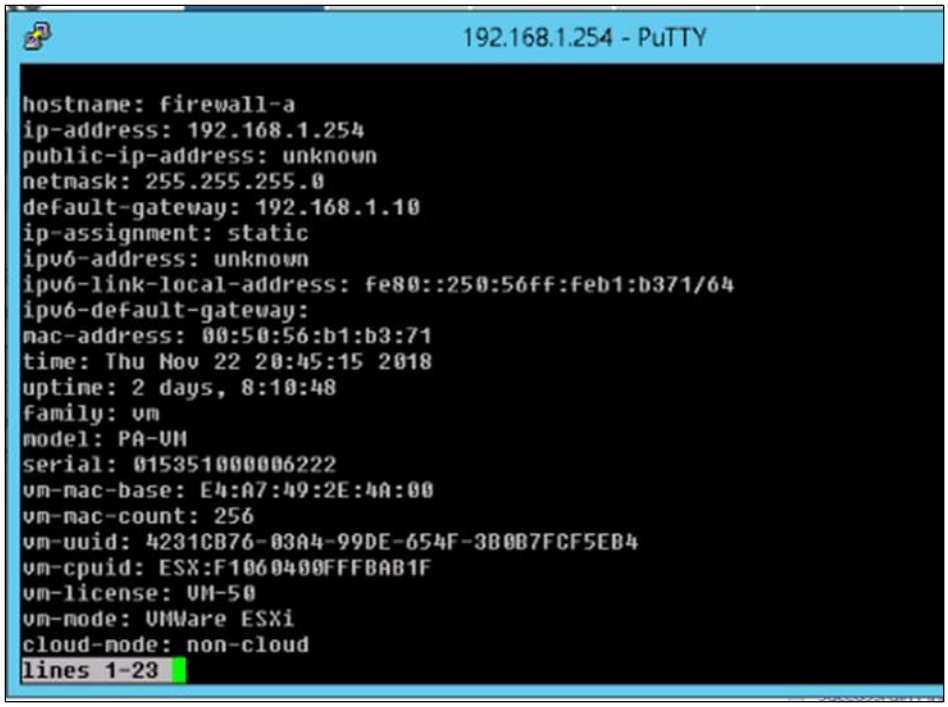

Panorama: Panorama is a Palo Alto Networks product that provides centralized web-based management, reporting, and logging for multiple firewalls. Use Panorama for centralized policy and firewall management to increase operational efficiency in managing and maintaining a distributed network of firewalls. If six or more firewalls are deployed in your network, you should use Panorama to reduce the complexity and administrative overhead needed to manage configuration, policies, software, and dynamic content updates. The Panorama web interface is similar to the firewall web interface, but with additional management functions.

XML API: The XML API provides a representational state transfer (REST)-based interface to access firewall configurations, operational status, reports, and packet captures from the firewall. An API browser is available on the firewall at https://<firewall>/api, where <firewall> is the hostname or IP address of the firewall. You can use this API to access and manage your firewall through a third-party service, application, or script.

The PAN-OS XML API can be used to automate tasks such as:

- Create, update, and modify firewall and Panorama configurations
- Execute operational mode commands, such as restarting the system or validating configurations
- Retrieve reports
- Manage users through User-ID
- Update dynamic objects without having to modify or commit new configurations

Interface Management Profiles

Management of Palo Alto Networks firewalls is not limited to using a dedicated MGT interface or console port. Data interfaces on the data plane also can be used as management interfaces. If the MGT interface goes down, you can continue to manage the firewall by allowing management access over another data interface. Each data interface includes configurations for binding various services to them:

- HTTPS (default)
- SSH (default)
- Ping (default)
- Telnet
- HTTP
- SNMP
- Response Pages
- User-ID

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. 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.

HTTPS includes the web interface service and should be included on at least one data interface. The **Permitted IP Addresses** field allows an Access Control List to be included, thus restricting access to only specified IP addresses for any interface with this profile assigned. If no IP addresses are added to the list of **Permitted IP Addresses**, then any IP address is allowed. After at least one IP address is added to the list, only those IP addresses are allowed access.

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, the firewall denies management access for all IP addresses, protocols, and services by default.

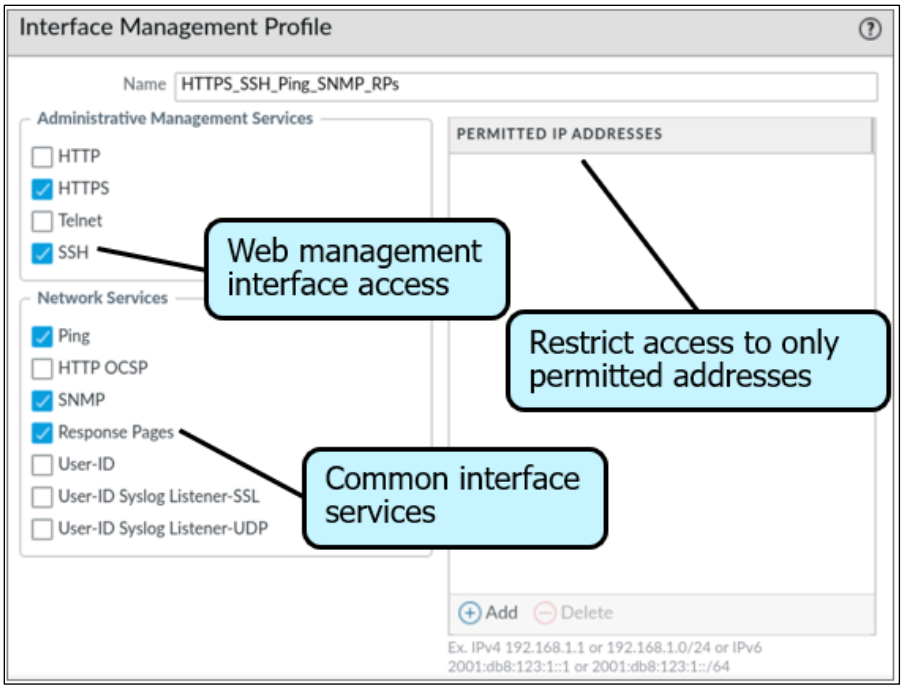
Interface Management Profile
Name
HTTPS_SSH_Ping_SNMP_RPs
Administrative Management Services
HTTP
HTTPS
Telnet
SSH
Network Services
Ping
HTTP OCSP
SNMP
Response Pages
User-ID
User-ID Syslog Listener-SSL
User-ID Syslog Listener-UDP
PERMITTED IP ADDRESSES
Add
Delete
Ex. IPv4 192.168.1.1 or 192.168.1.0/24 or IPv6 2001:db8:123:1::1 or 2001:db8:123:1::/64
Web management interface access
Restrict access to only permitted addresses
Common interface services

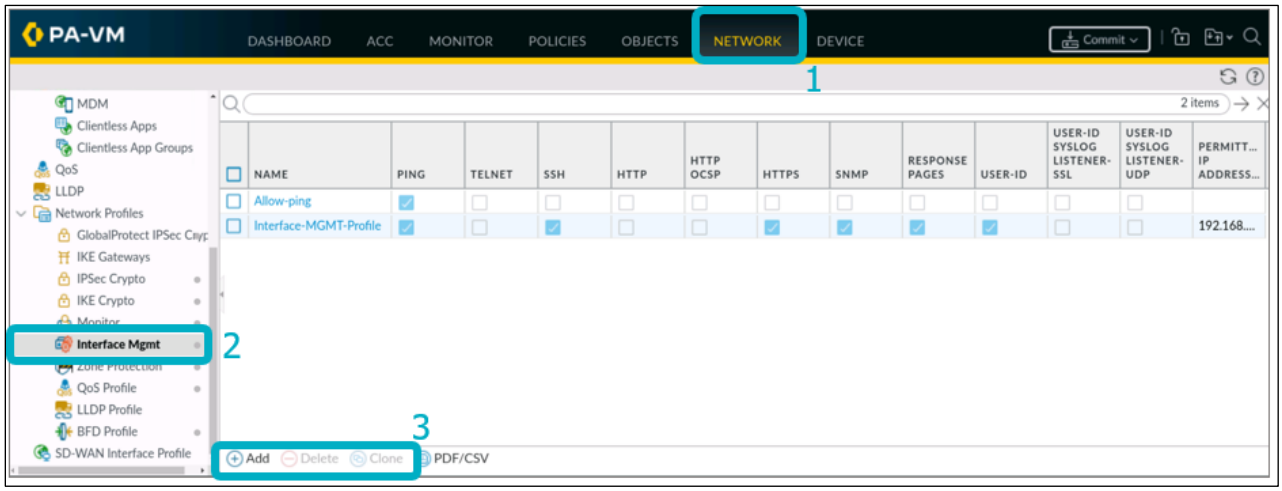
PA-VM
DASHBOARD
ACC
MONITOR
POLICIES
OBJECTS
NETWORK
DEVICE
Commit
1
MDM
Clientless Apps
Clientless App Groups
QoS
LLDP
Network Profiles
IKE Gateways
IPSec Crypto
IKE Crypto
Interface Mgmt
QoS Profile
LLDP Profile
BFD Profile
SD-WAN Interface Profile
2
2 items
NAME
PING
TELNET
SSH
HTTP
HTTP OCSP
HTTPS
SNMP
RESPONSE PAGES
USER-ID
USER-ID SYSLOG LISTENER-SSL
USER-ID SYSLOG LISTENER-UDP
PERMITT... IP ADDRESS...
Allow-ping
Interface-MGMT-Profile
192.168....
3
Add
Delete
Clone
PDF/CSV

Firewall Web Interface: Dashboard

Firewall Dashboard

The firewall **Dashboard** provides information in a condensed format. It is the main screen for web interface management.

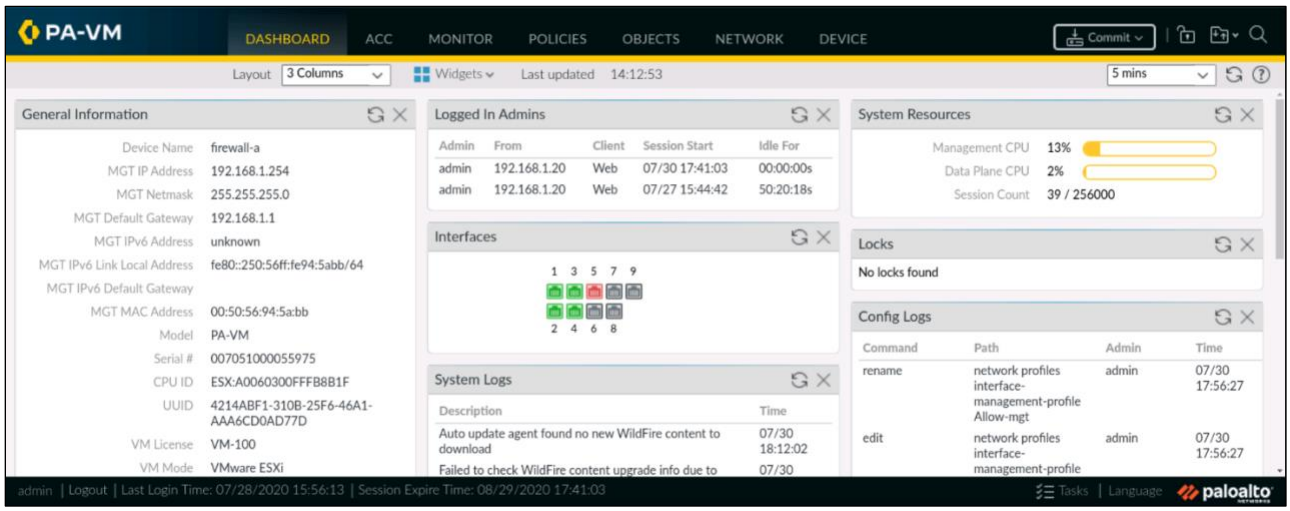

The **Dashboard** is customizable and allows you to determine which widgets to display:

- Application widgets:
    - ACC Risk Factor
    - Top Applications
    - Top High Risk Applications
- Logs widgets:
    - Config Logs
    - Data Filtering Logs
    - System Logs
    - Threat Logs
    - URL Filtering Logs
- System widgets:
    - General Information
    - High Availability

- Interfaces
- Locks
- Logged In Admins
- System Resources

Functional Category Tabs

Management of the firewall is conducted using seven category tabs, which are listed and briefly described as follows:

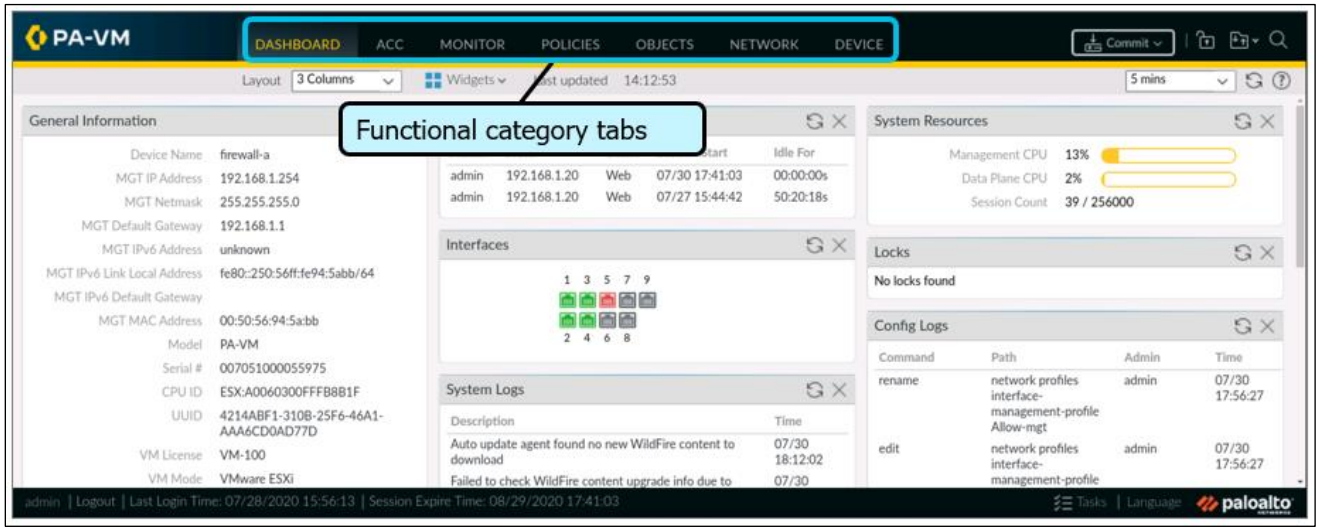

- **Dashboard**: Provides general information such as device name, MGT IP address, and licensing information. This page can be augmented by adding widgets.
- **ACC**: Uses the firewall logs to graphically depict traffic trends on your network
- **Monitor**: Provides logging visibility and the ability to run packet captures
- **Policies**: Allows the creation of policies such as Security policy and NAT policy
- **Objects**: Allows the creation of objects such as Address objects
- **Network**: Allows the configuration of network parameters such as interfaces and zones
- **Device**: Allows the configuration of system information such as the hostname or certificates

Tasks Icon

The **Tasks** icon appears at the bottom right. Select it to display the tasks that you, other administrators, or the PAN-OS software has initiated since the last firewall reboot (for example, manual commits or automatic FQDN refreshes):

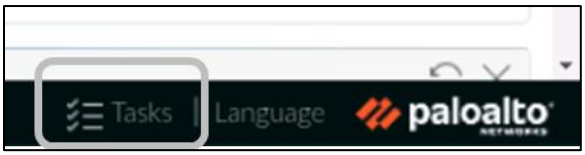

Service Routes

By default, the firewall uses the management interface to communicate with various servers including those for External Dynamic Lists (EDLs), DNS, email, and Palo Alto Networks updates servers. The management interface also is used to communicate with Panorama. Service routes are used so that the communication between the firewall and servers goes through the data ports on the data plane. These data ports require appropriate Security policy rules before external servers can be accessed.

Configuring Service Routes

Go to **Device > Setup > Services > Service Route Configuration > Customize** and configure the appropriate service routes. See the following figure:

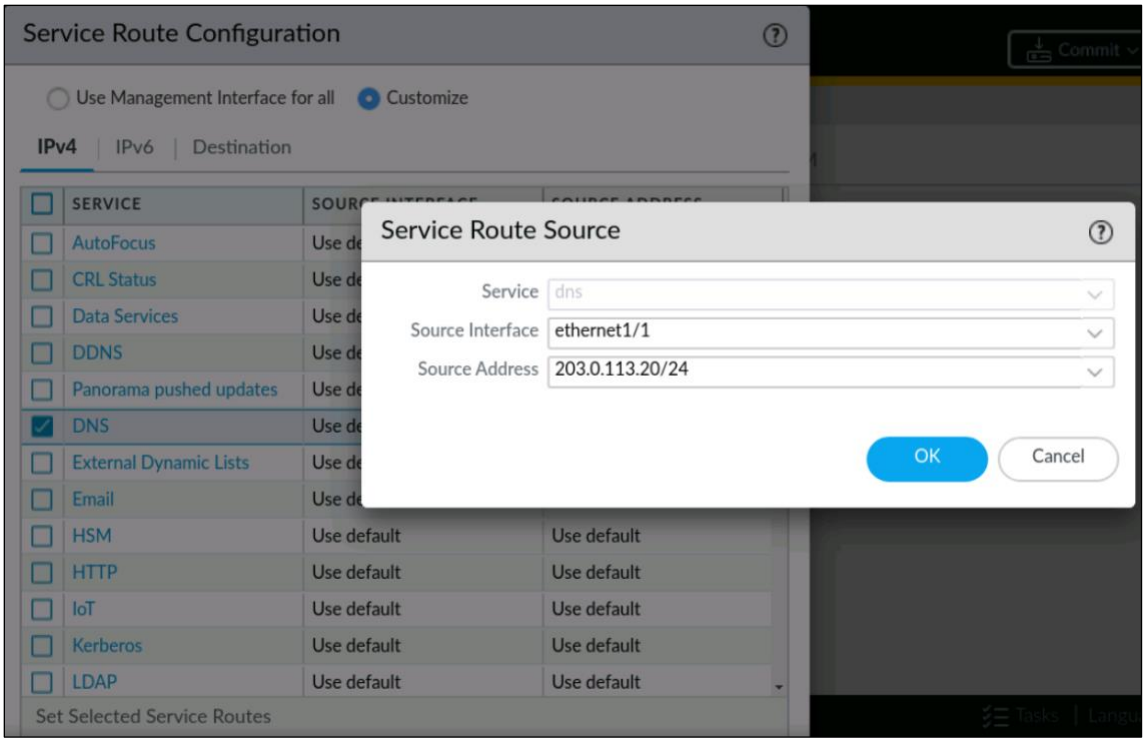

To configure service routes for non-predefined services, you can manually enter the destination addresses on the **Destination** tab:

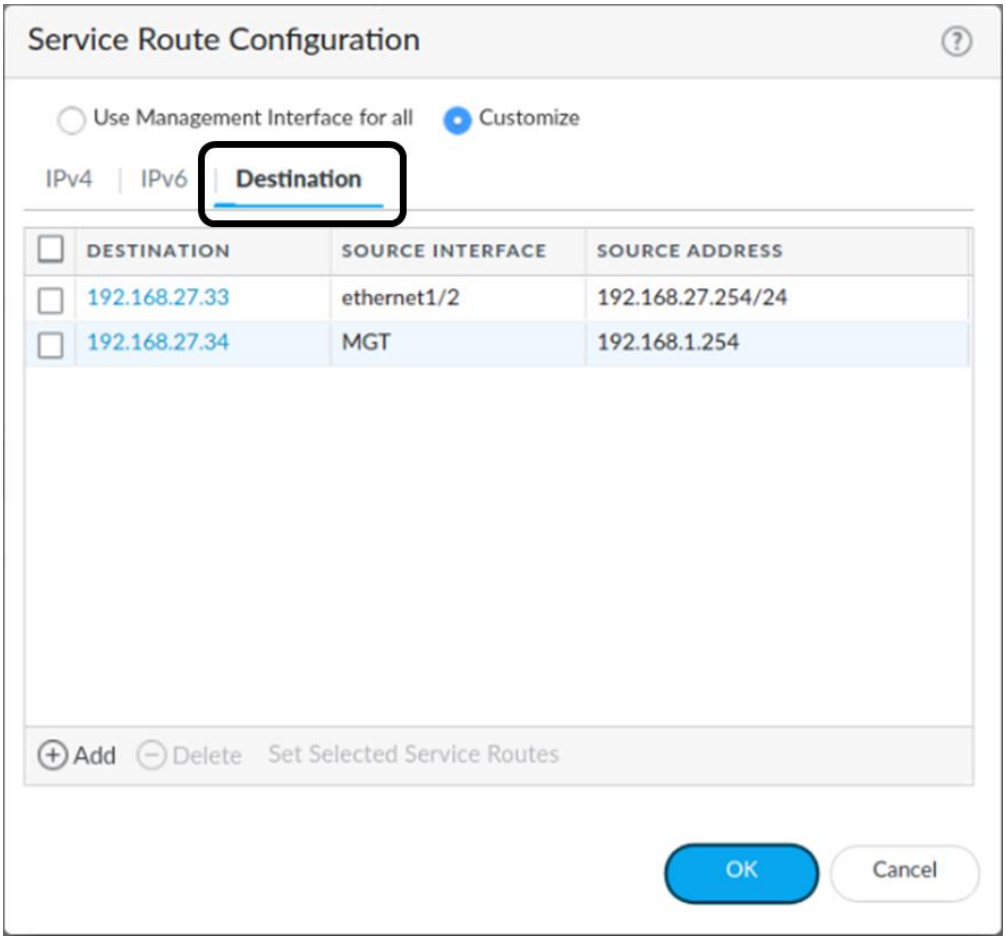

In the example shown, the service route for 192.168.27.33 is configured to source from the data plane's ethernet1/2 interface, which has a source IP address of 192.168.27.254.

Firewall Services

Palo Alto Networks firewalls integrate with three important services: DNS, DHCP, and NTP. DNS and NTP must be set during the initial firewall configuration.

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. You must configure your firewall with at least one DNS server so it can resolve hostnames.

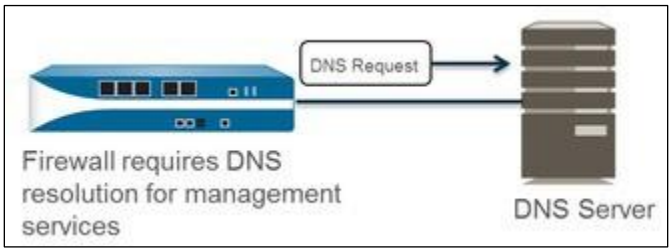

Configuring DNS

To configure DNS, select **Device > Setup > Services > Services_gear_icon**. On the **Services** tab, for DNS, click **Servers** and enter the **Primary DNS Server** address and **Secondary DNS Server** address. Click **OK** and **Commit**:

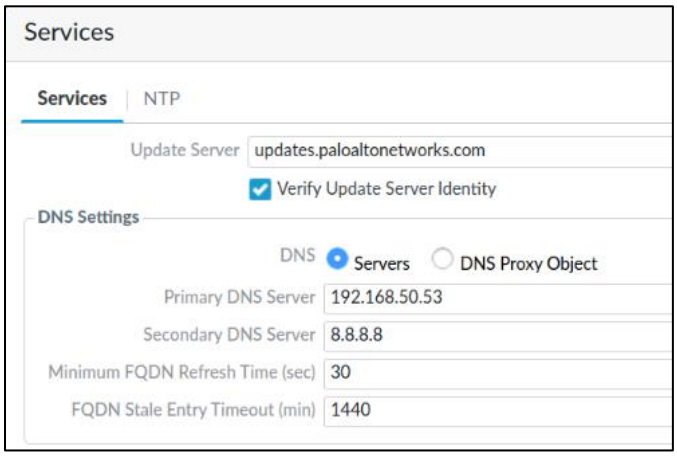

DHCP

A Palo Alto Networks firewall acting as a DHCP client (host) can request an IP address and other configuration settings from a DHCP server. The use of DHCP saves time and effort because users need not know the network's addressing plan or other options, such as default gateway, they are inheriting from the DHCP server.

Configuration parameters that DHCP can learn dynamically include:

- IP address for MGT port
- Netmask
- Default gateway
- At least one DNS server address

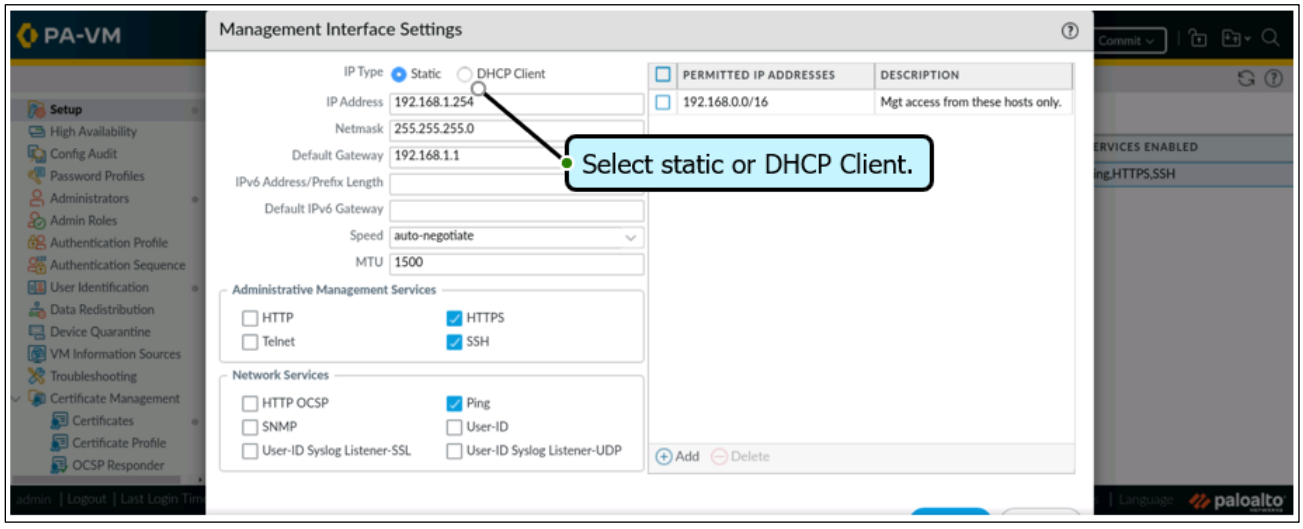

NTP

NTP client information is optional but recommended. The NTP information can be obtained via DHCP if the firewall is configured as a DHCP client.

Configuring NTP

Select **Device > Setup > Services > Services_gear_icon**:

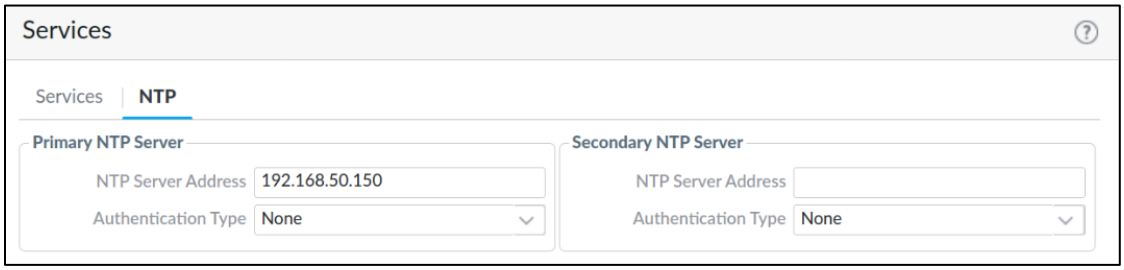

Sample Questions

Q1. What are two firewall management methods? (Choose two.)

a) CLI

b) RDP

c) VPN

d) XML API

Q2. Which two devices are used to connect a computer to the firewall for management purposes? (Choose two.)

a) rollover cable

b) serial cable

c) RJ-45 Ethernet cable

d) USB cable

Q3. What is the default IP address on the MGT interfaces of a Palo Alto Networks firewall?

a) 192.168.1.1

b) 192.168.1.254

c) 10.0.0.1

d) 10.0.0.254

Q4. What are the two default services that are available on the MGT interface? (Choose two.)

a) HTTPS

b) SSH

c) HTTP

d) Telnet

Q5. True or false. Service route traffic has Security policy rules applied against it.

a) true

b) false

Q6. Service routes may be used to forward which two traffic types out a data port? (Choose two.)

a) External Dynamic Lists

b) MineMeld

c) Skype

d) Palo Alto Networks updates

2.2 Provision local administrators and assigning role-based authentication

There are times when you may want to configure the firewall and Panorama to provide local authentication for administrators and end users. For example, you may require special user accounts that you don't manage through the directory servers that your organization reserves for regular accounts. Another example is when you 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 these cases, you can use the following authentication methods:

- **Local database authentication**: Below are the high-level process steps to configure local database authentication:

    1. To add a user account to the local database, 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**. **Enable** the account.

- **Local authentication without a database**: Below are the high-level process steps to configure firewall administrative accounts or Panorama administrative accounts without creating a database of users and user groups:

    1. To add an administrative account on the firewall, select **Device** > **Administrators** and **Add** an account.
    2. Enter a user **Name** for the administrator.
    3. Select an Authentication Profile. If the firewall uses local authentication without a local user database, select **None** and enter a **Password**.
    4. Select the **Administrative Type**.
    5. Select a **Password Profile**.
    6. Click **OK** and **Commit**.

Reference

Administrative Role Types
https://docs.paloaltonetworks.com/pan-os/9-0/pan-os-admin/firewall-administration/manage-firewall-administrators/administrative-role-types.html#id8b324bf1-eac8-40e1-82d5-6f82ff761fa9

2.3 Define firewall configurations

Manage Configurations Using Candidate and Running Configurations

All configuration changes in a Palo Alto Networks firewall are done to a *candidate configuration*, which resides in memory on the control plane. A commit activates the changes since the last commit and installs the running configuration on the data plane, where it will become the *running configuration*.

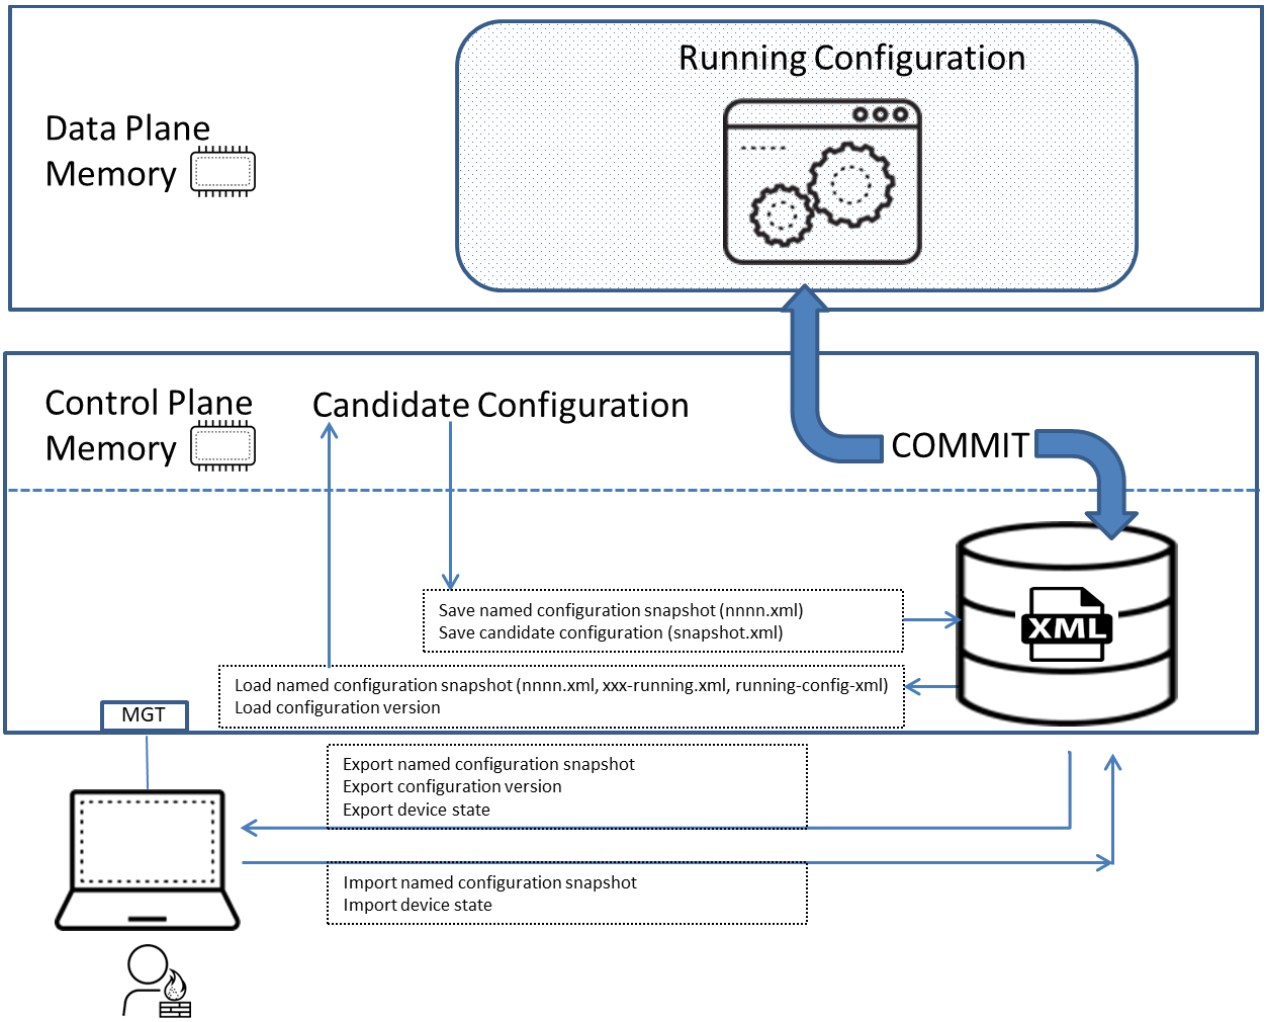

Candidate Configuration

The act of saving your changes to the candidate configuration does not activate those changes. A commit must be performed on the firewall to activate the changes and to cause the candidate configuration to become the running configuration. The commit can be done either via the web interface or the CLI.

The candidate configuration can be saved as either a default snapshot file (snapshot.xml) or as a custom-named snapshot file (<custom_name>.xml). However, a firewall does not automatically save the candidate configuration to persistent storage; you must manually save the candidate configuration. If the firewall reboots before you commit your changes, you can revert the candidate configuration to the current snapshot to restore changes you made between the last commit and the last snapshot using the **Revert to last saved configuration** option.

Running Configuration

The running configuration is a configuration that is saved within a file named running-config.xml. The running configuration exists in data-plane memory, where it is used to control firewall traffic and operate the firewall. A commit operation is necessary to write the candidate configuration to the running configuration.

After you commit changes, the firewall automatically saves a new *version* of the running configuration that is timestamped. You can load a previous version of the running configuration using the **Load configuration version** option. 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 commits that the firewall initiates automatically, such as FQDN refreshes.

If a system event or administrator action causes a firewall to reboot, the firewall automatically reverts to the current version of the running configuration.

Manage Running and Candidate Configurations

Palo Alto Networks firewall configurations are managed using five categories, which are found under **Device > Setup > Operations** and are described in the next sections:

- Revert
- Save
- Load
- Export
- Import

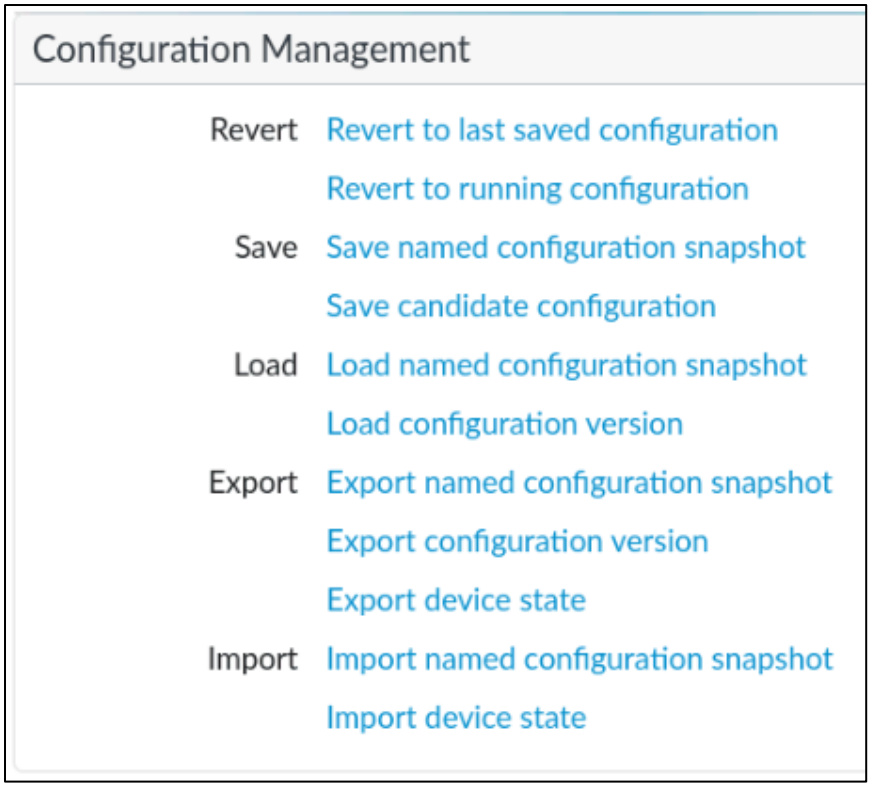

As a best practice, periodically save candidate configurations.

Revert to Last Saved Configuration

This option restores the default snapshot (snapshot.xml) of the candidate configuration (the snapshot that you create or overwrite when you click **Device > Setup > Operations > Save candidate configuration** or **Save** at the top right of the web interface). This option restores the last saved candidate configuration from the local drive. The current candidate configuration is overwritten. This quick restore is useful when you work on "hot" boxes.

The first message asks if you want to continue with the revert:

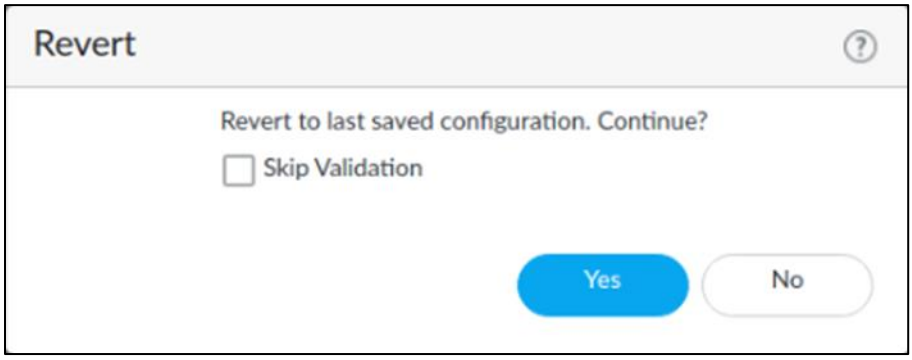

The second message informs you which file has been reverted:

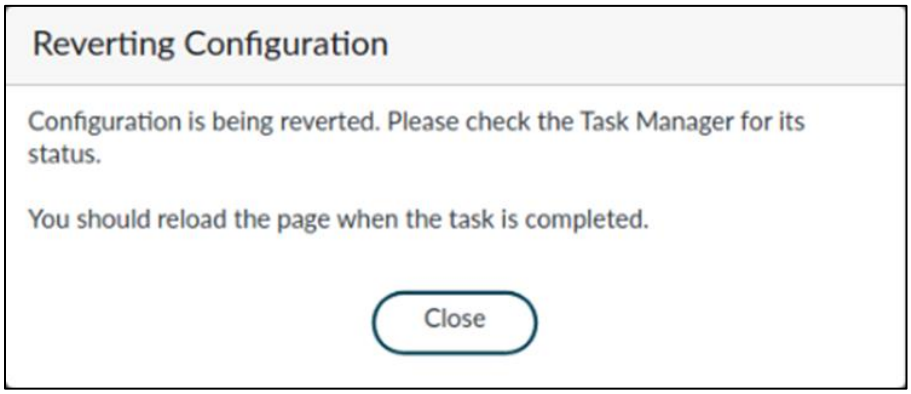

Revert to Running Configuration

This option restores the current running configuration. This operation undoes all the changes you made to the candidate configuration since the last commit and restores the config from the running-config.xml file.

The first message asks if you want to continue with the revert:

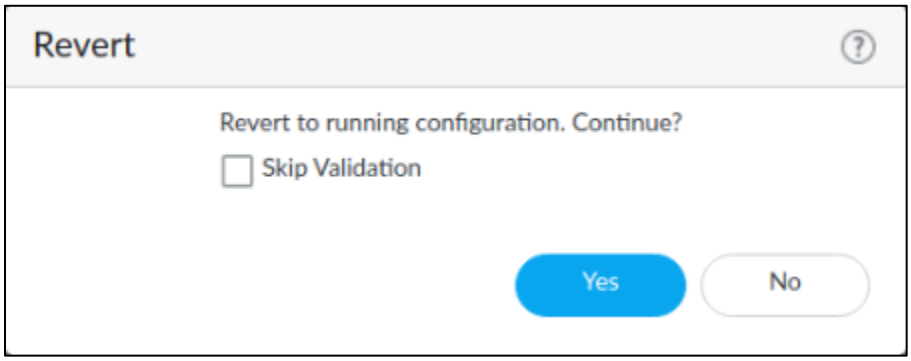

The second message informs you the firewall is being reverted.

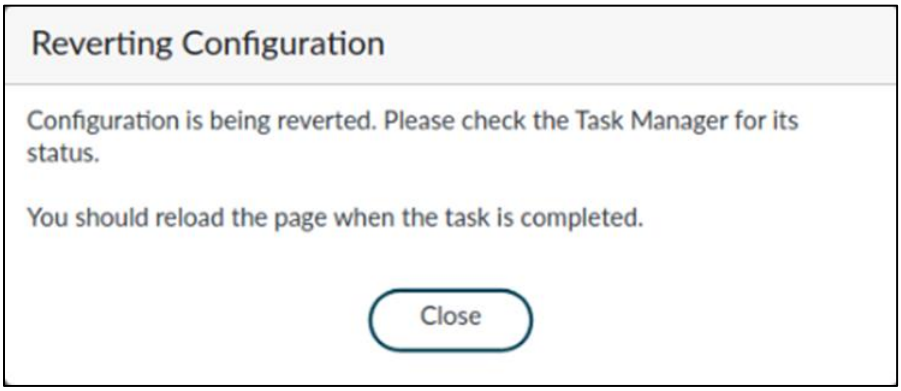

Save Named Configuration Snapshot

This option creates a candidate configuration snapshot that does not overwrite the default snapshot (snapshot.xml). Enter a custom name for the snapshot or select an existing snapshot to overwrite. This function is useful when you create a backup file or a test configuration file that could be downloaded for a further modification or testing in the lab environment.

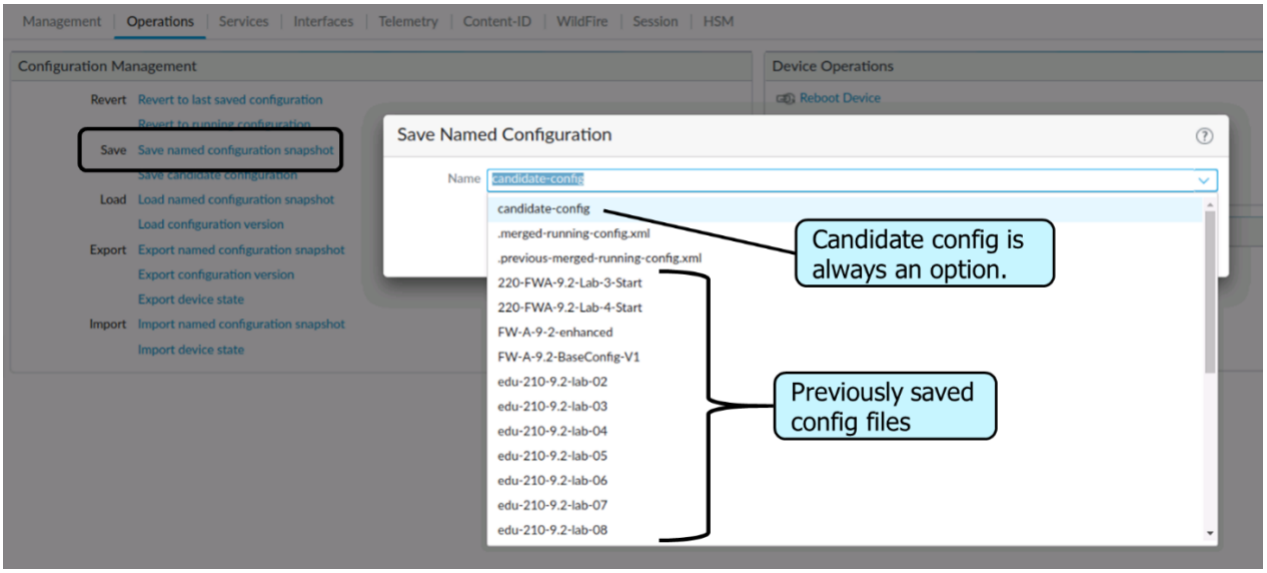

Save Candidate Configuration

This option creates or overwrites the default snapshot (snapshot.xml) of the candidate configuration (the snapshot that you create or overwrite when you click **Device > Setup > Operations > Save candidate configuration** or **Save** at the top right of the web interface).

Load Named Configuration Snapshot

This option overwrites the current candidate configuration with one of the following:

- Custom-named candidate configuration snapshot (instead of the default snapshot)
- Custom-named running configuration that you imported
- Current running configuration (running-config.xml)

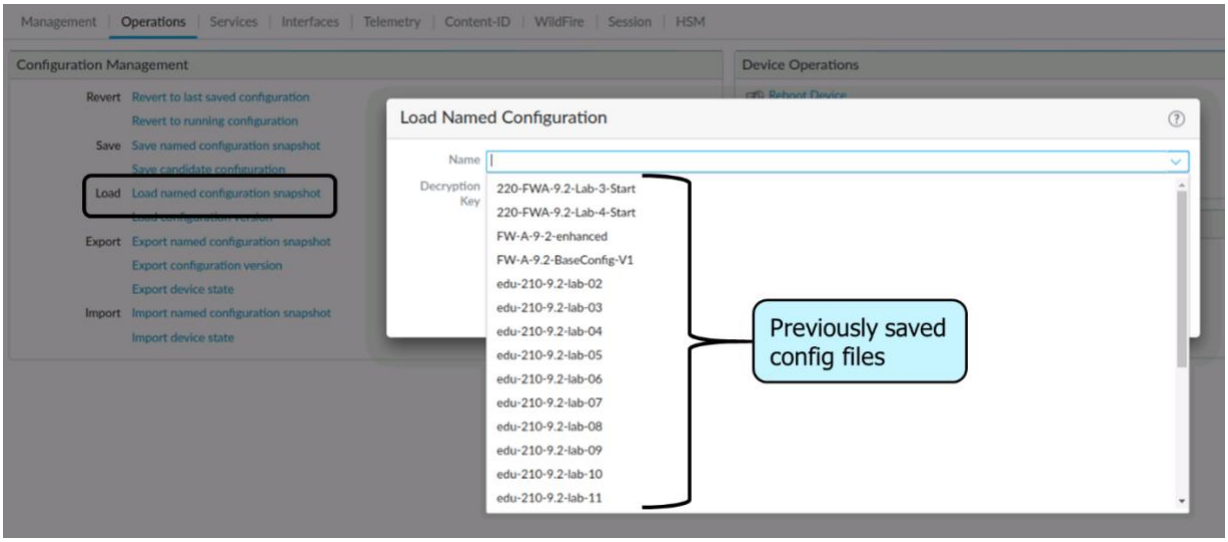

Load Configuration Version

This option overwrites the current candidate configuration with a previous version of the running configuration that is stored on the firewall. The firewall creates a timestamped version of the running configuration whenever a commit is made.

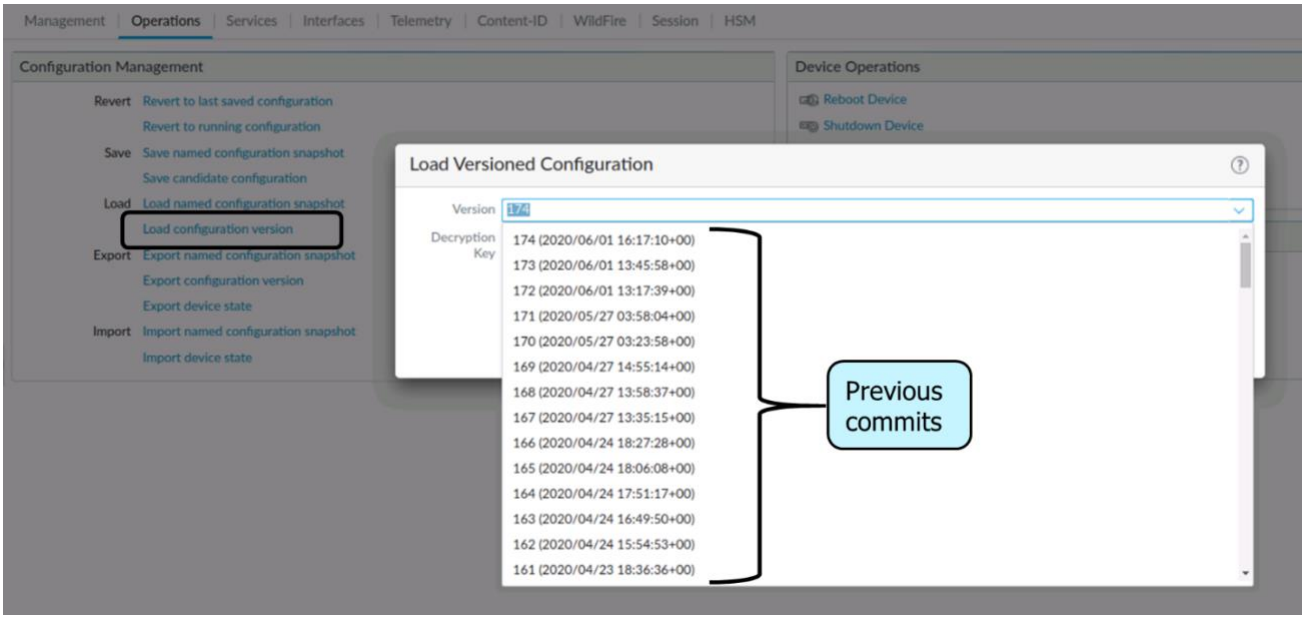

Export Named Configuration Snapshot

This option exports the current running configuration, a candidate configuration snapshot, or a previously imported configuration (candidate or running). The firewall exports the configuration as an XML file with the specified name. You can save the snapshot in any network location. These exports often are used as backups. These XML files also can be used as templates for building other firewall configurations.

Export Configuration Version

This option exports a version of the running configuration as an XML file.

Export Device State

This option exports the firewall state information as a file. In addition to the running configuration, the state information includes device group and template settings pushed from Panorama, if applicable. If the firewall is a GlobalProtect portal, the bundle also includes certificate information, a list of satellites that the portal manages, 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.

Import Named Configuration Snapshot

This option imports a running or candidate configuration as an XML file from any network location such as a host computer. Click **Browse** and select the configuration file to be imported. The XML file then can be loaded as a candidate configuration and even ultimately loaded as the running configuration if required.

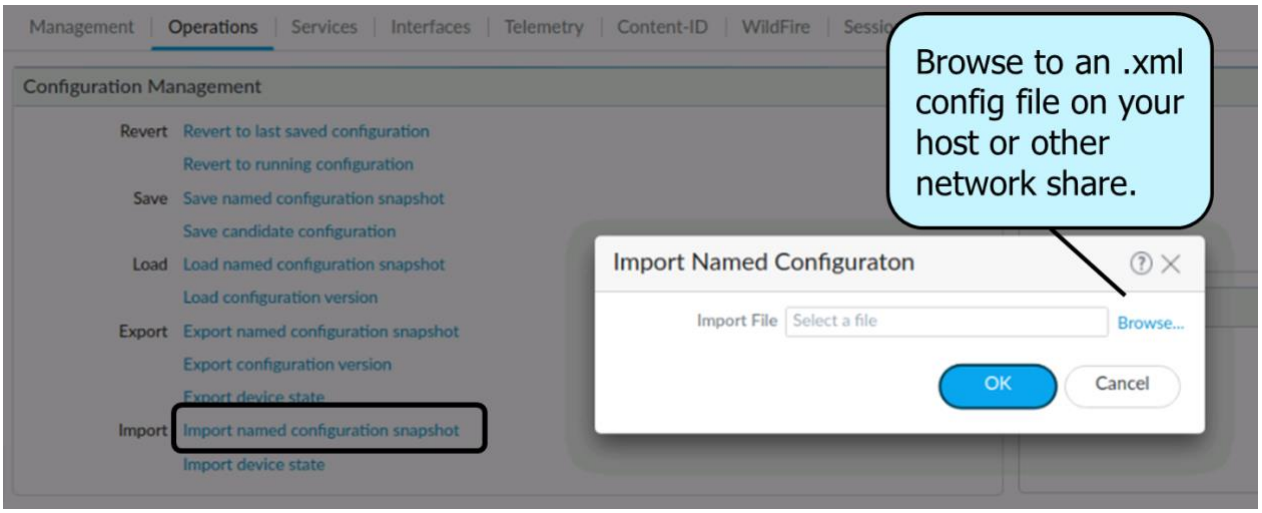

Import Device State

This option imports the state information file that you exported from a firewall using the **Export device state** option. The state information includes the running configuration and device group and template settings pushed from Panorama, if applicable. If the firewall is a GlobalProtect portal, the bundle also includes certificate information, a list of satellites, and satellite authentication information. If you replace a firewall or portal, you can restore the information on the replacement by importing the state bundle.

Sample Questions

Q1. Which firewall plane does the running configuration reside on?

a) management

b) control

c) data

d) security

Q2. Which firewall plane does the candidate configuration reside on?

a) management

b) control

c) data

d) security

Q3. Candidate config and running config files are saved as which file type?

a) TXT

b) HTML

c) XML

d) RAR

Q4. Which command must be performed on the firewall to activate any changes?

a) commit

b) save

c) load

d) import

Q5. Which command backs up configuration files to a remote network device?

a) import

b) load

c) copy

d) export

Q6. The command **`load named configuration snapshot`** overwrites the current candidate configuration with which three items? (Choose three.)

a) custom-named candidate configuration snapshot (instead of the default snapshot)

b) custom-named running configuration that you imported

c) snapshot.xml

d) current running configuration (running-config.xml)

e) Palo Alto Networks updates

2.4 Understand how to push policy updates to Panorama-managed firewalls

You can use Panorama to manage your firewalls. You'll need to enable the connection from the firewall to Panorama; to enable this connection, add a firewall as a managed device.

To add a firewall as a managed device, perform the following high-level process tasks:

1. Configure the firewall so it's accessible with Panorama over the network.
2. Configure each data interface on the firewall you plan to use and attach it to a **security zone**. This will allow you to push configuration and policy updates.
3. Add the **Panorama IP address** to the firewall.
4. Add one or more firewalls to Panorama (**Panorama > Managed Devices > Summary**).
5. Enter the firewall **Serial** number. Select the **Associate Devices** check box.
6. Assign the **Device Group**, **Template Stack**, **Collector Group**, and **Log Collector**.
7. Enable **Auto Push on 1st connect**.
8. Add a **tag**.
9. Select **Commit > Commit to Panorama** and **Commit**.

These are the high-level steps to start receiving push policy updates. If you've configured this correctly, you now have the option to automatically push the configuration to your newly added firewall when the firewall first connects to Panorama. This ensures that firewalls are immediately configured and ready to secure your networks.

References

- Set Up Zero Touch Provisioning https://docs.paloaltonetworks.com/panorama/9-1/panorama-admin/manage-firewalls/set-up-zero-touch-provisioning.html

- Transition a Firewall to Panorama Management
  https://docs.paloaltonetworks.com/panorama/9-1/panorama-admin/manage-firewalls/transition-a-firewall-to-panorama-management.html

2.5 Identify the types of dynamic updates and their purpose

Dynamic Updates

To ensure that you always are protected from the latest threats (including those that have not yet been discovered), you must keep your firewalls up-to-date with the latest content and software updates published by Palo Alto Networks. Palo Alto Networks regularly posts updates for application detection, threat protection, and GlobalProtect data files through dynamic updates.

The following diagram illustrates how often updated information is made available to the firewall:

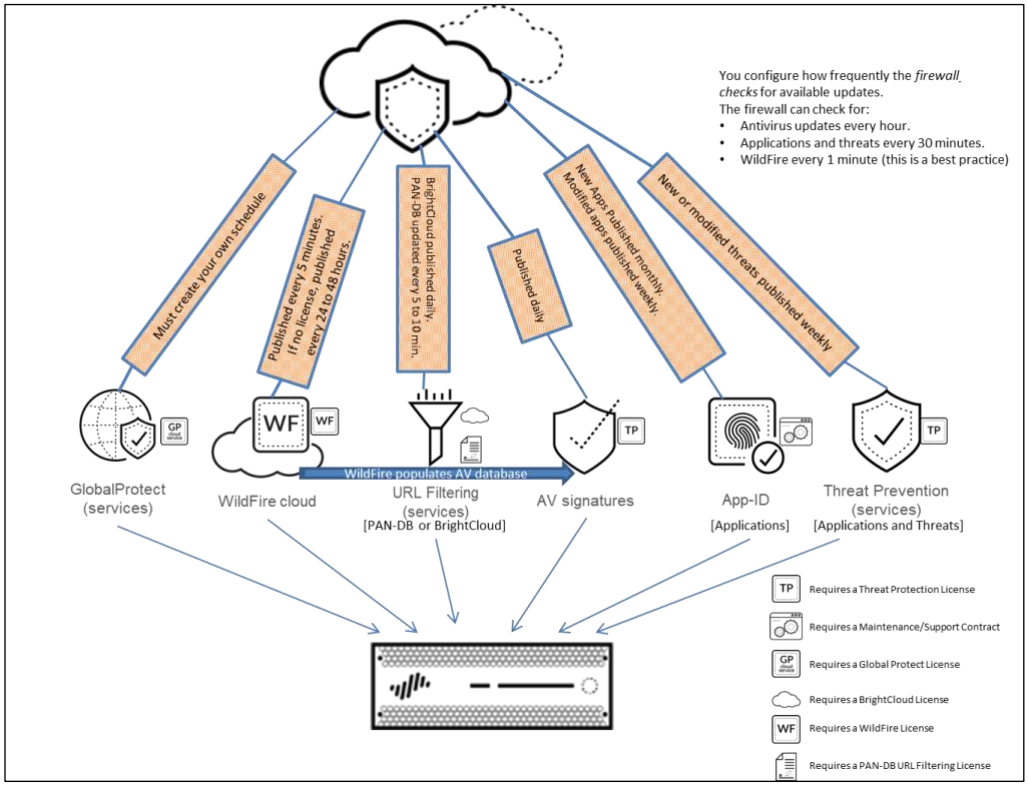

The following content updates are available, depending on which subscriptions you have:

- **Antivirus:** Includes new and updated antivirus signatures, including WildFire signatures and automatically generated command-and-control (C2) signatures. WildFire signatures detect malware seen first by firewalls from around the world. 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 applications are published monthly, and modified applications are published weekly. To best deploy application updates to ensure application availability, be sure to follow the best practices for Applications and Threats 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 and modified threat signatures and modified applications signatures are published weekly; new application signatures are published monthly. The firewall can retrieve the latest update within 30 minutes of availability. To best deploy application and threat updates based on your security and application availability needs, be sure to follow the best practices for Applications and Threats content updates.

- **GlobalProtect Data File:** Contains vendor-specific information for defining and evaluating host information profile (HIP) data returned by GlobalProtect clients. You must have a GlobalProtect license (subscription) and create an update schedule 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 to receive these updates and enable clientless VPN to function.

- **Palo Alto Networks (PAN-DB) URL filtering:** Complements 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. If URL filtering is enabled, all web traffic is compared against the URL filtering database, which contains a listing of millions of websites that have been categorized into 60 to 80 categories.

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. PAN-DB URL filtering requires a PAN-DB URL Filtering license.

Every 5 to 10 minutes a new version is published that contains updated categorization data and an incremented version number. Each time the Palo Alto Networks firewall sends a request to the cloud, it checks the current version number. If the number is different, the firewall upgrades the device's version to the current cloud version. The primary purpose of the frequency of updates is to leverage native integration with WildFire, which creates new signatures and records malicious URLs every 5 minutes.

- **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. End-of-sale was January 1, 2018 and end-of-support is July 21, 2021.
- **WildFire:** This update is available with a WildFire subscription and provides 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 WildFire updates every minute. If you have a Threat Prevention subscription and not a WildFire subscription, you must wait 24 to 48 hours for the WildFire signatures to be added into the antivirus update.
- **WF-Private:** Provides malware signatures generated by an on-premises WildFire appliance.

Downloading and Installing Updates

You can view the latest updates, read the release notes for each update, and then select the update you want to download and install. You also can revert to a previously installed version of an update.

Always review content Release Notes for the list of the newly identified and modified applications and threat signatures that the content release introduces:

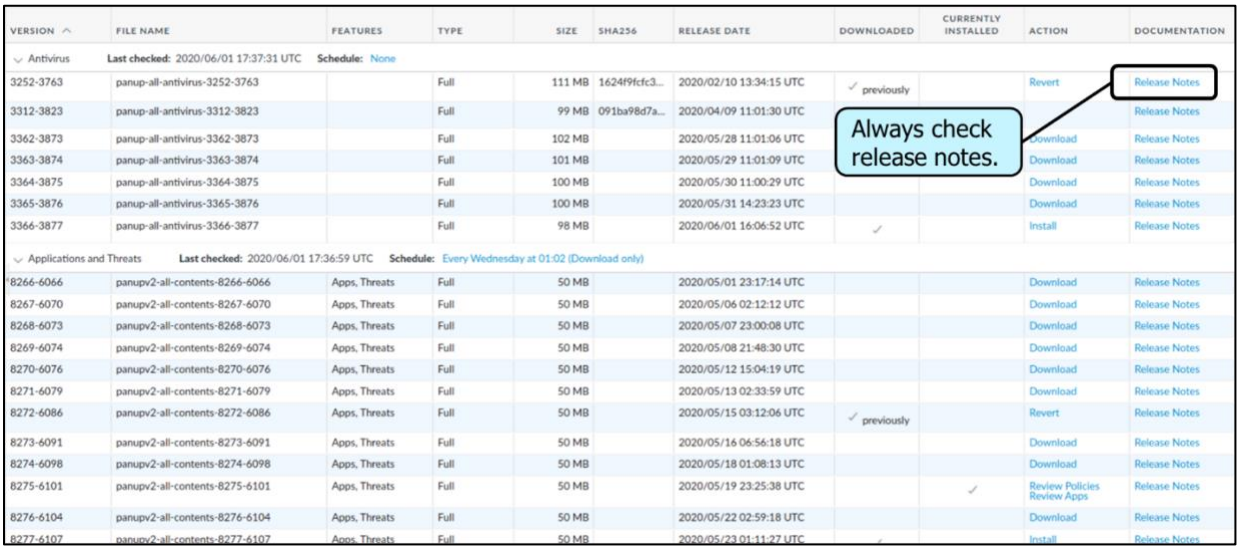

VERSION	FILE NAME	FEATURES	TYPE	SIZE	SHA256	RELEASE DATE	DOWNLOADED	CURRENTLY INSTALLED	ACTION	DOCUMENTATION
Antivirus	Last checked: 2020/06/01 17:37:31 UTC	Schedule: None								
3252-3763	panup-all-antivirus-3252-3763		Full	111 MB	1624f9fcfc3...	2020/02/10 13:34:15 UTC	✓ previously		Revert	Release Notes
3312-3823	panup-all-antivirus-3312-3823		Full	99 MB	091ba98d7a...	2020/04/09 11:01:30 UTC				Release Notes
3362-3873	panup-all-antivirus-3362-3873		Full	102 MB		2020/05/28 11:01:06 UTC			Download	Release Notes
3363-3874	panup-all-antivirus-3363-3874		Full	101 MB		2020/05/29 11:01:09 UTC			Download	Release Notes
3364-3875	panup-all-antivirus-3364-3875		Full	100 MB		2020/05/30 11:00:29 UTC			Download	Release Notes
3365-3876	panup-all-antivirus-3365-3876		Full	100 MB		2020/05/31 14:23:23 UTC			Download	Release Notes
3366-3877	panup-all-antivirus-3366-3877		Full	98 MB		2020/06/01 16:06:52 UTC	✓		Install	Release Notes
Applications and Threats	Last checked: 2020/06/01 17:36:59 UTC	Schedule: Every Wednesday at 01:02 (Download only)								
8266-6066	panupv2-all-contents-8266-6066	Apps, Threats	Full	50 MB		2020/05/01 23:17:14 UTC			Download	Release Notes
8267-6070	panupv2-all-contents-8267-6070	Apps, Threats	Full	50 MB		2020/05/06 02:12:12 UTC			Download	Release Notes
8268-6073	panupv2-all-contents-8268-6073	Apps, Threats	Full	50 MB		2020/05/07 23:00:08 UTC			Download	Release Notes
8269-6074	panupv2-all-contents-8269-6074	Apps, Threats	Full	50 MB		2020/05/08 21:48:30 UTC			Download	Release Notes
8270-6076	panupv2-all-contents-8270-6076	Apps, Threats	Full	50 MB		2020/05/12 15:04:19 UTC			Download	Release Notes
8271-6079	panupv2-all-contents-8271-6079	Apps, Threats	Full	50 MB		2020/05/13 02:33:59 UTC			Download	Release Notes
8272-6086	panupv2-all-contents-8272-6086	Apps, Threats	Full	50 MB		2020/05/15 03:12:06 UTC	✓ previously		Revert	Release Notes
8273-6091	panupv2-all-contents-8273-6091	Apps, Threats	Full	50 MB		2020/05/16 06:56:18 UTC			Download	Release Notes
8274-6098	panupv2-all-contents-8274-6098	Apps, Threats	Full	50 MB		2020/05/18 01:08:13 UTC			Download	Release Notes
8275-6101	panupv2-all-contents-8275-6101	Apps, Threats	Full	50 MB		2020/05/19 23:25:38 UTC		✓	Review Policies Review Apps	Release Notes
8276-6104	panupv2-all-contents-8276-6104	Apps, Threats	Full	50 MB		2020/05/22 02:59:18 UTC			Download	Release Notes
8277-6107	panupv2-all-contents-8277-6107	Apps, Threats	Full	50 MB		2020/05/23 01:11:27 UTC			Install	Release Notes

You can download updates directly from the Palo Alto Networks update server. You also can download the updates to another system such as a user desktop or a Panorama management appliance and then upload them to the firewall. Whether you download an update through the web or upload an update

from Panorama, the update will appear in the list of available updates at **Device > Dynamic Updates**. Click **Install** to install the updates.

Downloading Updates

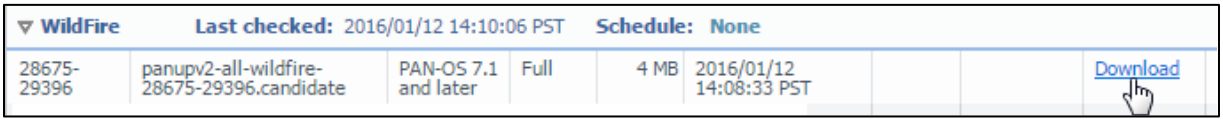

Installing Updates

Software Updates

PAN-OS updates are managed in the **Device > Software** section of the web interface. A final system reboot must be performed to put the new PAN-OS software into production. This reboot is disruptive and should be done during a change control window.

The software downloads are done over the MGT interface by default. A data interface can be used to download the software using a service route. The latest version of applications and threats must be installed to complete the software installation. If your firewall does not have internet access from the management port, you can download the software image from the Palo Alto Networks Support Portal and then manually **Upload** it to your firewall.

Before you upgrade to a newer version of software:

- Always review the release notes to determine any impact of upgrading to a newer version of software
- Ensure the firewall is connected to a reliable power source. A loss of power during an upgrade can make the firewall unusable.
- Although the firewall automatically creates a configuration backup, follow best practice and create and externally store a backup before you upgrade.

To upgrade to a newer version of software, complete the following steps:

1. Ensure you follow the correct upgrade path. When you upgrade, typically you must download the *x*.0 base release before you install the maintenance or feature release. For example, to upgrade from 7.*x*.*y* to 8.*x*.*y*, download both 8.0 and 8.*x*.*y*. 8.0 automatically is installed when you install 8.*x*.*y*.
2. Select **Device Software** and click **Check Now** to display the latest PAN-OS updates.
3. Locate and **Download** the applicable PAN-OS software.

4. After you download the image (or, for a manual upgrade, after you upload the image), **Install** the image:

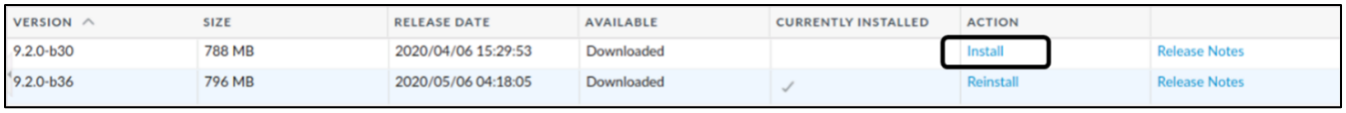

VERSION	SIZE	RELEASE DATE	AVAILABLE	CURRENTLY INSTALLED	ACTION	
9.2.0-b30	788 MB	2020/04/06 15:29:53	Downloaded		Install	Release Notes
9.2.0-b36	796 MB	2020/05/06 04:18:05	Downloaded	✓	Reinstall	Release Notes

5. After the installation completes successfully, reboot the firewall.

Sample Questions

Q1. True or false. A Palo Alto Networks firewall automatically provides a backup of the configuration during a software upgrade.

a) true

b) false

Q2. If you have a Threat Prevention subscription but not a WildFire subscription, how long must you wait for the WildFire signatures to be added into the antivirus update?

a) 1 to 2 hours

b) 2 to 4 hours

c) 10 to 12 hours

d) 24 to 48 hours

Q3. Which three actions should you complete before you upgrade to a newer version of software? (Choose three.)

a) Review the release notes to determine any impact of upgrading to a newer version of software.

b) Ensure the firewall is connected to a reliable power source.

c) Export the device state.

d) Create and externally store a backup before you upgrade.

e) Put the firewall in maintenance mode.

Identify the impact of deploying dynamic updates

Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to Palo Alto Networks firewalls. The firewalls access the web resources in the CDN to perform various App-ID and Content-ID functions. By default, the firewalls use the management port to access the CDN infrastructure for application updates, threat and antivirus signature updates, and access to the Palo Alto Networks WildFire® cloud. To ensure that you are always protected from the latest threats (including those that have not yet been discovered), you must keep your firewalls up to date with the latest content and software updates published by Palo Alto Networks.

App-ID updates have a special impact because new application definitions might affect current Security policy rules. PAN-OS software provides features to review the App-ID updates and modify the Security policy rules.

If your firewalls are managed by Panorama, the Panorama device can be the source of dynamic updates for managed firewalls and can configure the update schedule.

References

- Configure Content and Software Updates
  https://docs.paloaltonetworks.com/pan-os/10-0/pan-os-admin/software-and-content-updates/app-and-threat-content-updates/configure-app-threat-updates.html
- Manage New App-IDs Introduced in Content Releases
  https://docs.paloaltonetworks.com/pan-os/10-0/pan-os-admin/app-id/manage-new-app-ids-introduced-in-content-releases.html
- Managing dynamic updates from Panorama
  https://docs.paloaltonetworks.com/pan-os/10-0/pan-os-web-interface-help/panorama-web-interface/panorama-device-deployment/manage-software-and-content-updates.html

Sample Questions

Q4. After an Applications and Threats dynamic update is downloaded to the firewall, where can information about changes to the App-IDs be found?

a) Summary link in the log event detail reporting the dynamic update file download

b) Review Policies link at the bottom of the Security policy rules display

c) Review Apps link appearing next to the downloaded Applications and Threats file

d) Details link in the dynamic file availability announcement appearing in the News Feed widget on the dashboard

Q5. The GlobalProtect Data File dynamic update contains which kinds of data?

a) GlobalProtect client package software updates for Windows and Macintosh

b) list of available connection points for Prisma Access

c) HIP check detection data for the GlobalProtect clients

d) updates to cypher suites used by the GlobalProtect client

Q6. When application details are viewed in the App-ID database, which field indicates that a recent change to the application might affect your Security policy rules?

a) Name

b) Depends on

c) Previously Identified As

d) App-ID Enabled

2.6 Identify what a security zone is and how to use it

Security Zones

Palo Alto Networks firewalls use security zones to analyze, control, and log network traffic as it traverses from one zone interface to another zone interface. Zones logically group networks that contain particular types of traffic that are contained within defined security classifications. Examples of such classifications are Internet, Data Center Applications, Users, IT Infrastructure, and Customer Data.

Security zones are divided into two broad categories: Intrazone and Interzone. Security zones contain one or more physical or virtual interfaces. An interface can belong to only one zone. Intrazone traffic, by default, allows traffic to flow between interfaces that exist in the same zone. Interzone traffic, by default, denies traffic from flowing between interfaces that exist in different zones.

Security policy rules are applied to zones (not interfaces) to allow or deny traffic, apply QoS, perform NAT, apply security profiles, or set logging parameters. Security policy rules are described in another section of this study guide.

The following diagram is an example of network segments partitioned into multiple zones based on their security classification. The zones and the corresponding Security policy rules should be made as definitive as possible to reduce your network's attack surface. All zone names are custom names that are defined by the firewall administrator. There are five primary zone types (Tap, Virtual Wire, Layer 2, Layer 3, and Tunnel, that support only specific interface types, which also are depicted in the following diagram. Different zone and interface types can be used simultaneously on different physical firewall interfaces. Tunnel zones became available in PAN-OS 8.0 and are used for a feature named tunnel content inspection.

A sixth zone type named External is a special zone that is available only on some firewall models. The External zone allows traffic to pass between virtual systems when multiple virtual systems are configured on the same firewall. Virtual systems are supported only on the PA-2000, PA-3000, PA-4000, PA-5000, and PA-7000 Series firewalls. The External zone type is visible in the drop-down list only when it is supported by a firewall with the virtual systems feature enabled.

Note that MGT and HA interfaces are not assigned to a zone.

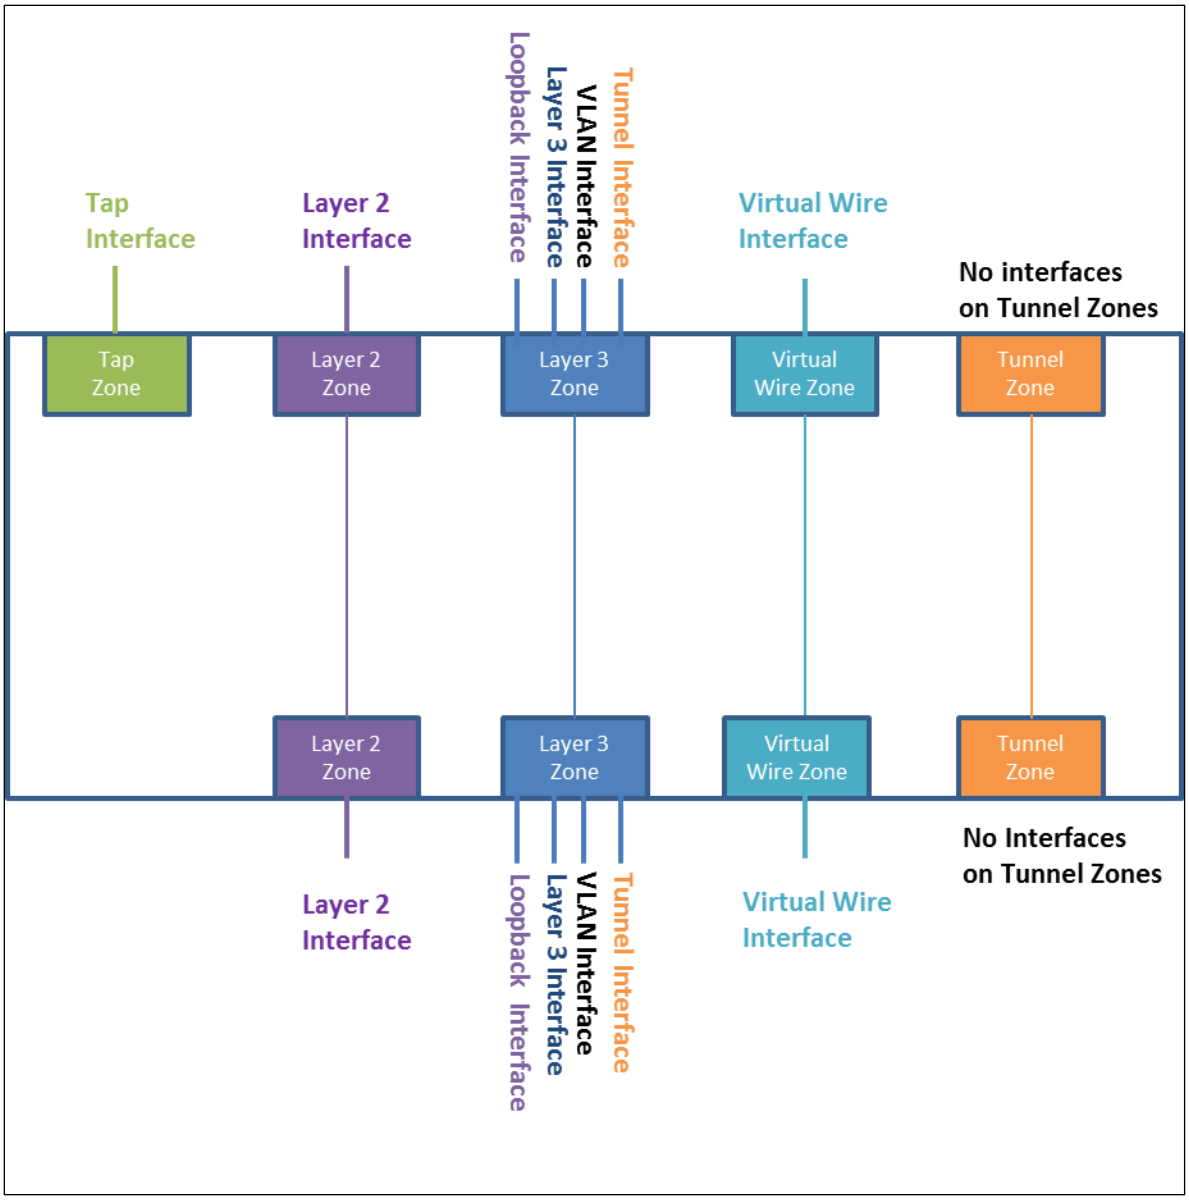

Zones need to be created and configured by assigning a zone name and specifying the zone type. Interfaces do not have to be configured prior to the zone's creation; they can be assigned to a zone later. Note that zone names are case-sensitive.

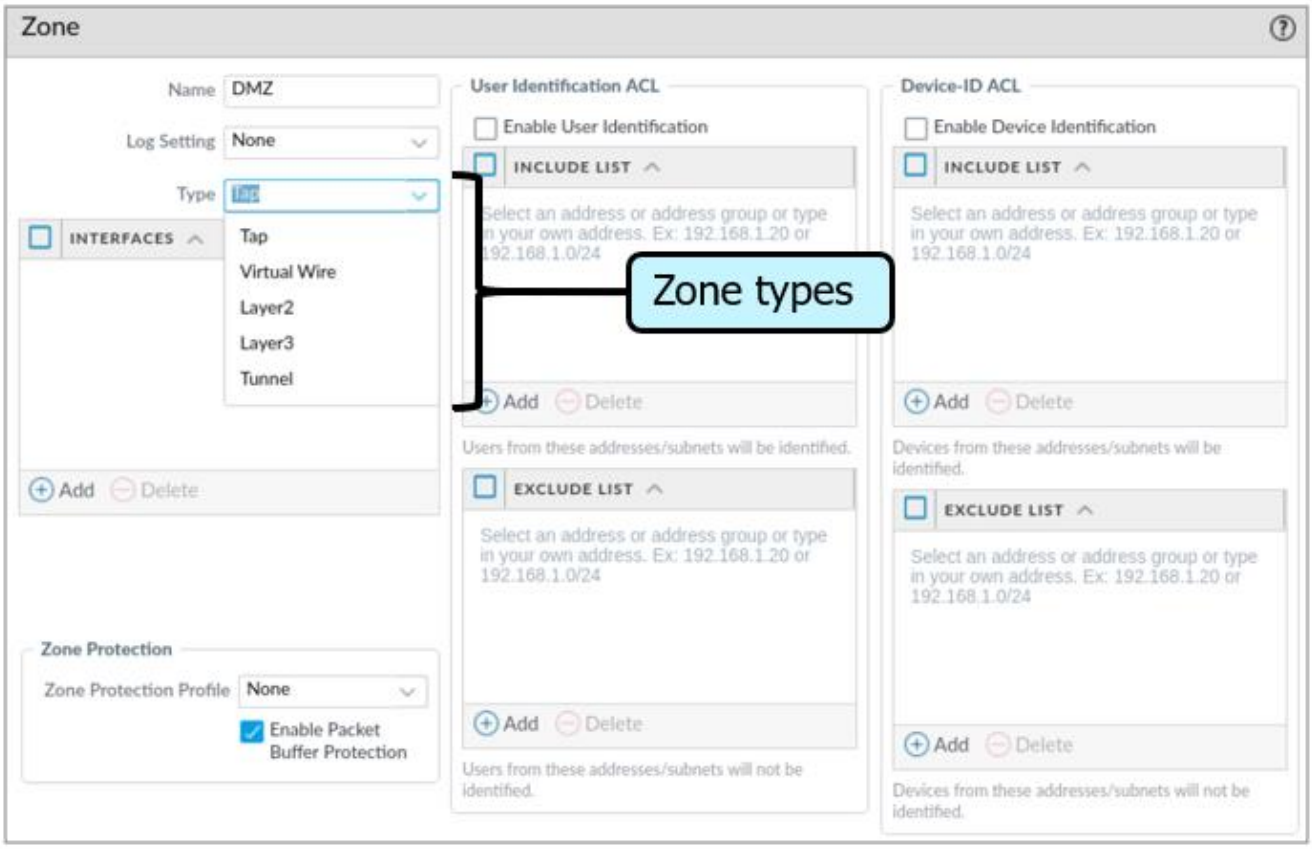

The following figure shows that the Layer 3 zone allows four interface types: Layer 3 (Ethernet1/4 and 1/5), loopback, sdwan, tunnel, and vlan:

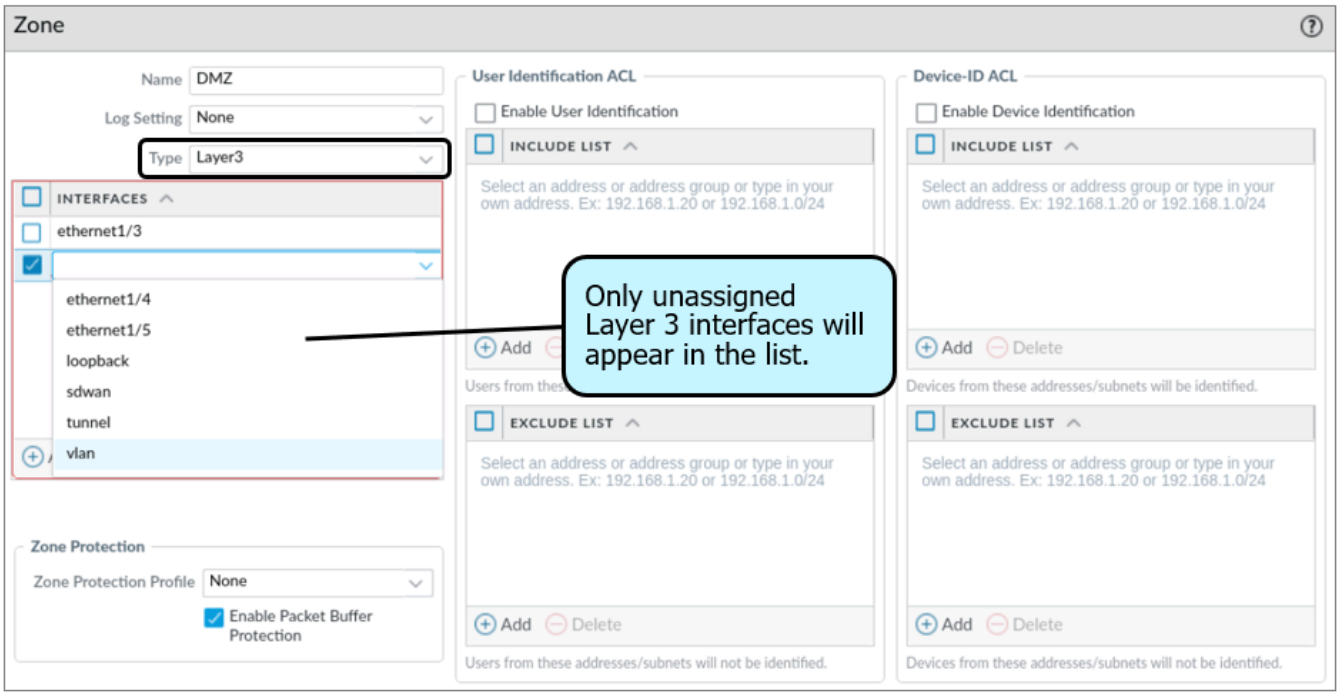

Sample Questions

Q1. Which two default zones are included with the PAN-OS software? (Choose two.)

a) Interzone

b) Extrazone

c) Intrazone

d) Extranet

Q2. Which two zone types are valid? (Choose two.)

a) trusted

b) tap

c) virtual wire

d) untrusted

e) dmz

Q3. Which two statements about interfaces are correct? (Choose two.)

a) Interfaces must be configured before you can create a zone.

b) Interfaces do not have to be configured before you can create a zone.

c) An interface can belong to only one zone.

d) An interface can belong to multiple zones.

Q4. Which two interface types can belong in a Layer 3 zone? (Choose two.)

a) Loopback

b) Tap

c) Tunnel

d) Virtual Wire

Q5. What are used to control traffic through zones?

a) access lists

b) Security policy lists

c) Security policy rules

d) Access policy rules

2.7 Identify and configure firewall interfaces

Types of Ethernet Interfaces

PAN-OS software has the following Ethernet interface types: Tap, Virtual Wire, Layer 2, Layer 3, and HA. (High Availability [HA] interfaces are not discussed in this section). A firewall can be configured with multiple instances of each interface type to accommodate its functional requirements within a network. The following figure shows how a firewall can be used in Tap, Virtual Wire, Layer 2, or Layer 3 mode.

Ethernet Interface Types

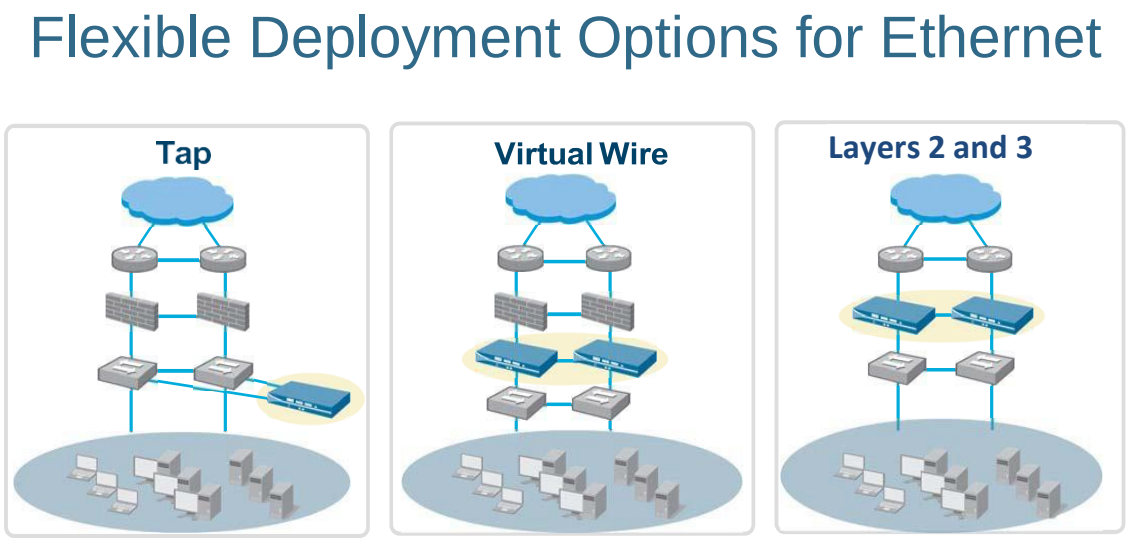

• App-ID, Content-ID, and User-ID visibility without inline deployment • Traffic logged to provide visibility	• SSL decryption (no encryption) • Allows NAT	• All the Virtual Wire mode capabilities including Layer 2 or Layer 3 services: virtual routers, VPN, and routing protocols

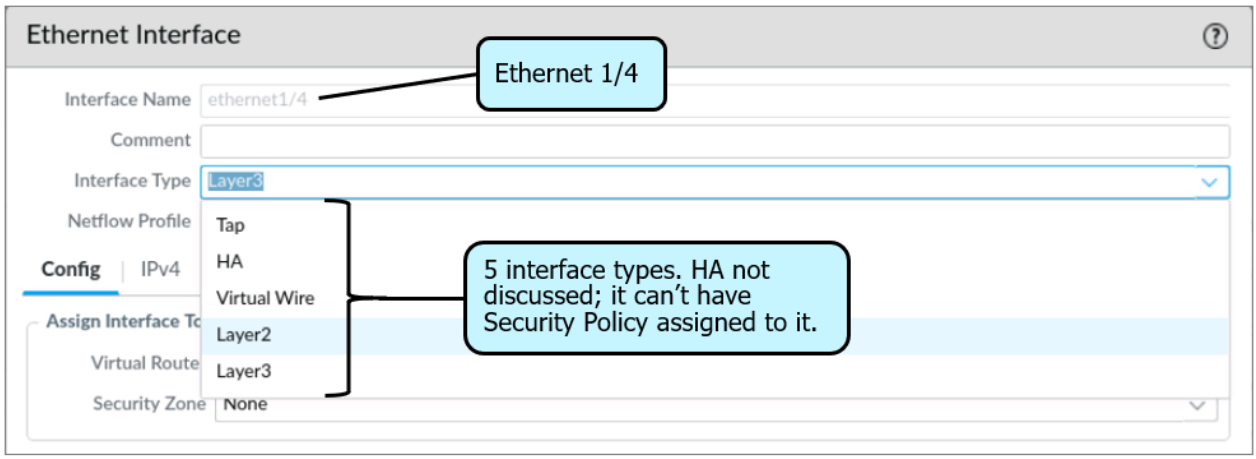

Other available interface types include the following:

- **Decrypt Mirror**: This feature enables decrypted traffic from a firewall to be copied and sent to a traffic collection tool that can receive raw packet captures, such as NetWitness or Solera, for archiving and analysis. Decrypt Mirror often is used to route decrypted traffic through an external interface to a data loss prevention (DLP) service. DLP is a product category for products that scan internet-bound traffic for keywords and patterns that identify sensitive information. Note that a free license is required to use this feature. This feature is not available on the VM-Series firewalls.

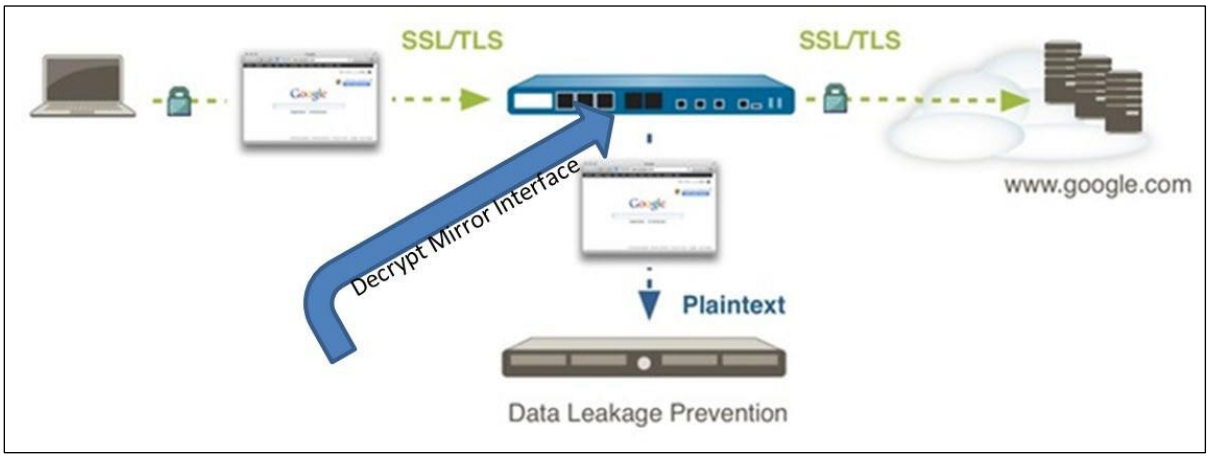

- **Log card:** For PA-7000 Series firewalls only. A log card data port performs log forwarding for syslog, email, Simple Network Management Protocol (SNMP), and WildFire file forwarding. One data port on a PA-7000 must be configured as a log card interface because the MGT interface cannot handle all the logged traffic.

- **Aggregate:** Used to bundle multiple physical HA3, Virtual Wire, Layer 2, or Layer 3 interfaces into a logical interface for better performance (via load balancing) and redundancy using IEEE 802.1AX (LACP) link aggregation. The interface types to be bundled must be the same. VM-Series models do not support Aggregate Ethernet (AE) interface groups.

- **HA interface:** Each HA interface has a specific function. One HA interface is for configuration synchronization and heartbeats; the other HA interface is for state synchronization. If active/active high availability is enabled, the firewall also can use a third HA interface to forward packets.
- **Management:** MGT interfaces are used to manage a firewall using a network cable.
- **Loopback:** Loopback interfaces are Layer 3 virtual interfaces that connect to virtual routers in the firewall. Loopback interfaces are used for multiple network engineering and implementation purposes. They can be destination configurations for DNS sinkholes, GlobalProtect service interfaces (portals and gateways), routing identification, and more.
- **Tunnel:** A tunnel interface is a logical (virtual) interface used with VPN tunnels to deliver encrypted traffic between two endpoints. The tunnel interface must belong to a security zone before policy can be applied, and it must be assigned to a virtual router to use the existing routing infrastructure. A tunnel interface does not require an IP address to route traffic between the sites. An IP address is required only if you want to enable tunnel monitoring or if you are using a dynamic routing protocol to route traffic across the tunnel.
- **SD-WAN:** Create and configure a virtual SD-WAN interface to specify one or more physical, SD-WAN-capable Ethernet interfaces that go to the same destination, such as to a specific hub or to the internet. In fact, all links in a virtual SD-WAN interface must be the same type: all VPN tunnel links or all direct internet access (DIA) links. An SD-WAN interface definitions works with an SD-WAN Interface Profile that defines the characteristics of the ISP connections. Details about these interfaces and their configuration are beyond the scope of the PCNSA certification.

Tap, Virtual Wire, Layer 2, and Layer 3 Interfaces

- **Tap**: A Tap interface monitors traffic that is connected to a network switch's MIRROR/SPAN port. This mirrored traffic is forwarded by a switch port to a firewall's Tap interface and is analyzed for App-ID, User-ID, Content-ID, and other traffic, just like any other normal data traffic that would pass through the firewall. Before traffic can be logged, a Security policy must be configured that includes the Tap zone. Tap interfaces are easy to deploy and can be implemented without disruption to your existing network. Tap mode offers visibility in the Traffic log and also in the ACC tab. The information can be used to help configure Security policy rules, and to make other firewall configuration changes. Tap traffic is not managed (blocked, allowed, or shaped). TAP interfaces must be assigned to a Tap zone.

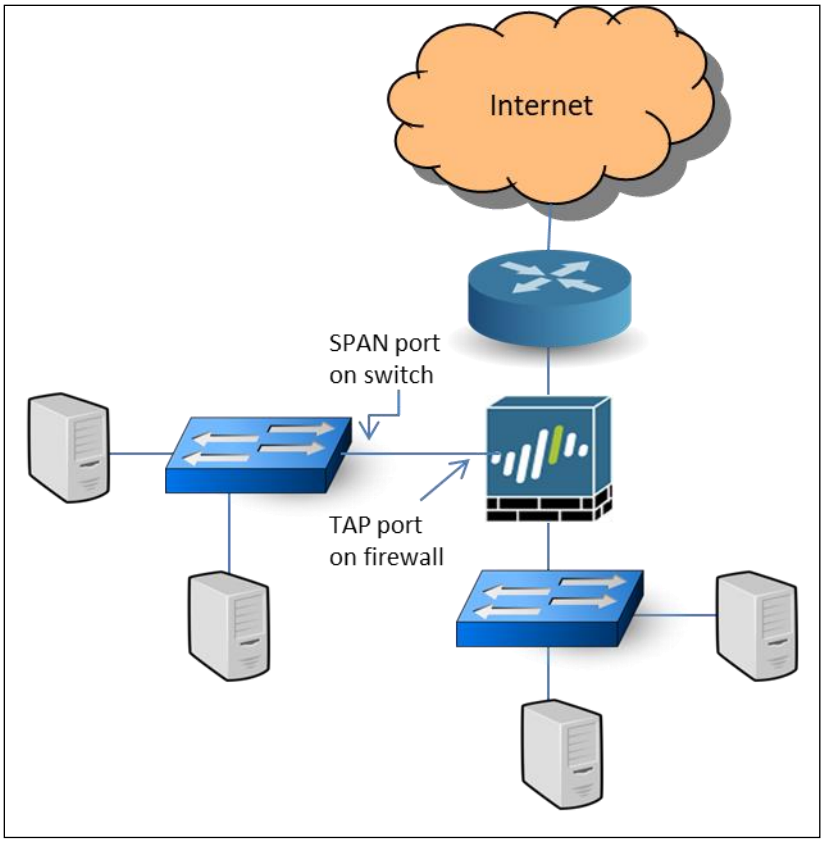

To configure a Tap interface, go to **Network > Interfaces > Ethernet > <select_interface>**.

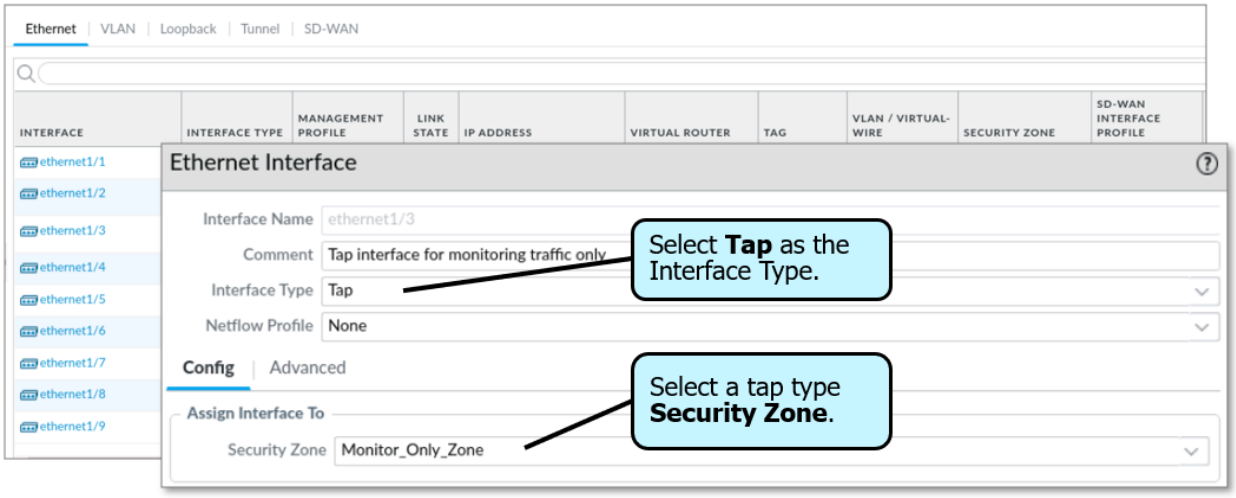

Virtual Wire

A Virtual Wire interface is used to pass traffic through a firewall by binding two Ethernet interfaces and allowing traffic to pass between them. Virtual Wire interfaces often are placed between an existing firewall and a secured network to enable analysis of the traffic before actually migrating from a legacy firewall to a Palo Alto Networks firewall.

- No IP or MAC addresses are assigned to Virtual Wire interfaces. No routing or switching is done on a Virtual Wire interface. A Virtual Wire interface that receives a frame or packet 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 Virtual Wire interface and on to the network device connected to it. A virtual wire requires no changes to adjacent network devices. A virtual wire can bind two Ethernet interfaces of the same medium (both either copper or fiber) or bind a copper interface to a fiber interface.
- Two Virtual Wire interfaces, each in a virtual wire zone (the zone can be the same or different), and a Virtual Wire object are required to complete a virtual wire configuration. The following figure shows one interface in one zone (Internet) and the other interface in another zone (Inside). If both interfaces are in different zones (interzone traffic), all traffic will be inspected by Security policy rules until sessions can be established, and then you can check for User-ID, App-ID, and Content-ID, and perform logging, QoS, decryption, LLDP, zone protection, DoS protection, and NAT.
- If both interfaces are in the same zone (intrazone traffic), all the traffic would be allowed by default, and sessions can be easily established. However, you also can check for User-ID, App-ID, and Content-ID, and perform logging, QoS, decryption, LLDP, zone protection, DoS protection, and NAT.
- Virtual Wire interfaces can be subdivided into Virtual Wire subinterfaces that can be used to classify traffic according to VLAN tags, IP addresses, IP ranges, or subnets. Use of subinterfaces enables you to separate traffic into different zones for more granular control than regular (non-subinterface) Virtual Wire interfaces.

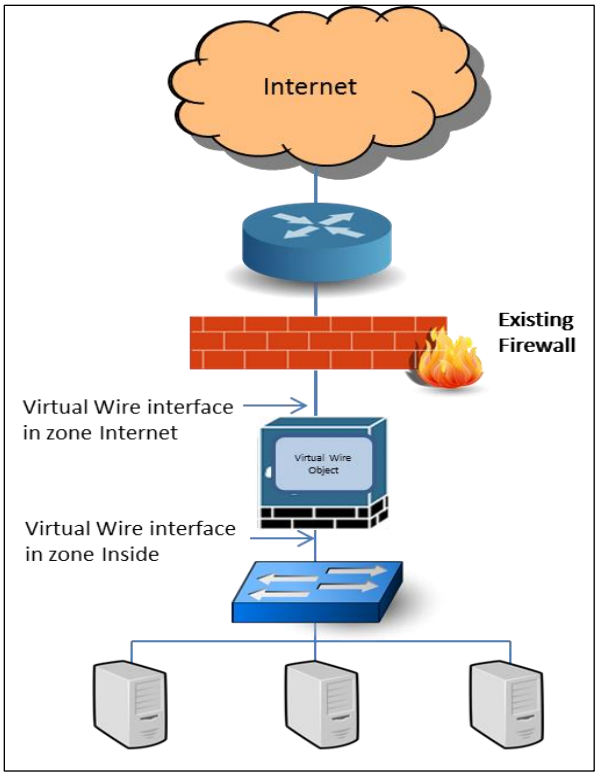

To configure a Virtual Wire object, go to **Network > Virtual Wires > Add**:

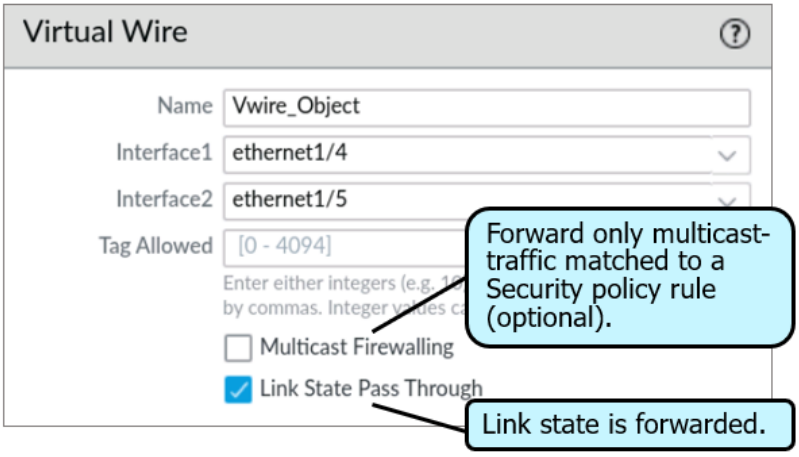

To configure a Virtual Wire interface, go to **Network > Interfaces > Ethernet > <select_interface>**:

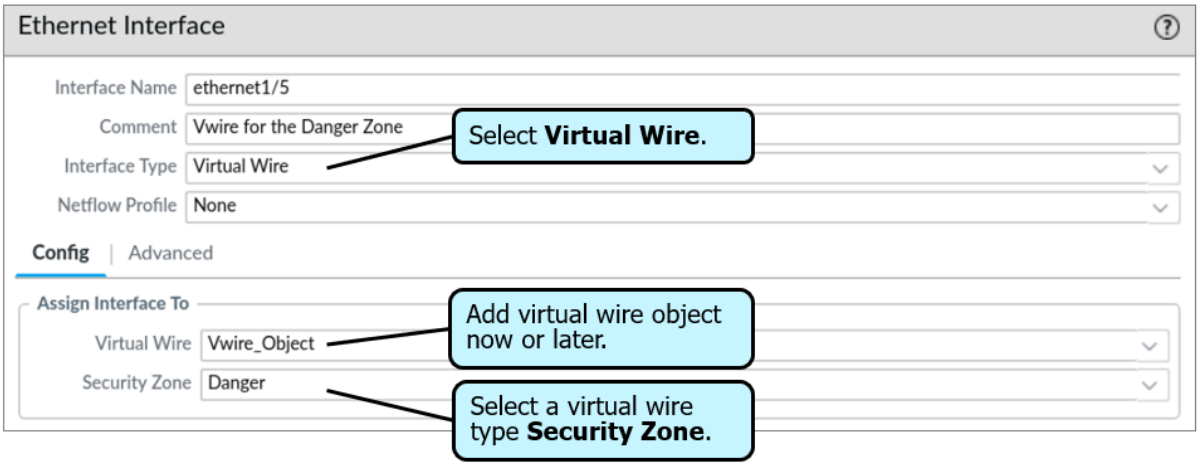

Virtual Wire Subinterfaces

Virtual wire deployments can use Virtual Wire subinterfaces to separate traffic into different zones. Virtual Wire subinterfaces provide flexibility in enforcing distinct policies when you need to manage traffic from multiple customer networks. Virtual Wire subinterfaces enable you to control and separate traffic by specifying criteria such as VLAN tags and IP classifiers. IP classifiers consist of host IP addresses, IP subnets, and IP ranges. Assign each subinterface to a different zone, and then enforce Security policy rules for the traffic that matches the defined criteria. Note that zones can belong to separate virtual systems.

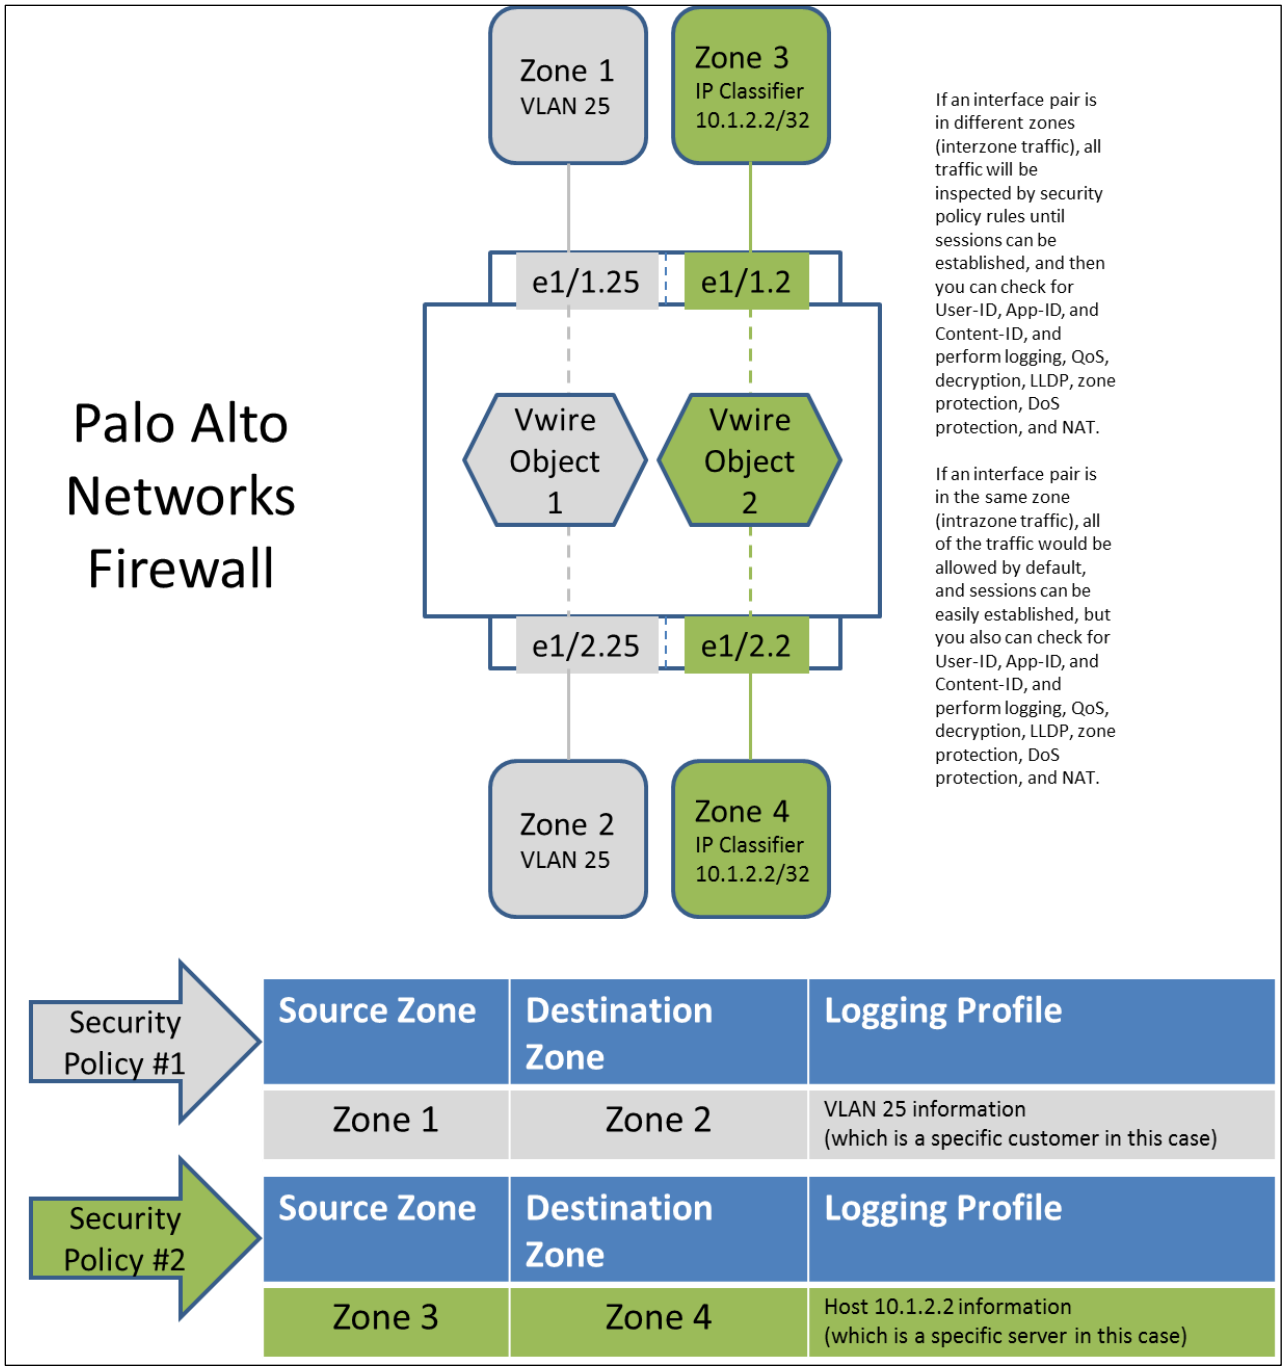
Zone 1
VLAN 25
Zone 3
IP Classifier
10.1.2.2/32
e1/1.25
e1/1.2
Vwire
Object
1
Vwire
Object
2
e1/2.25
e1/2.2
Zone 2
VLAN 25
Zone 4
IP Classifier
10.1.2.2/32
Palo Alto
Networks
Firewall
If an interface pair is in different zones (interzone traffic), all traffic will be inspected by security policy rules until sessions can be established, and then you can check for User-ID, App-ID, and Content-ID, and perform logging, QoS, decryption, LLDP, zone protection, DoS protection, and NAT.
If an interface pair is in the same zone (intrazone traffic), all of the traffic would be allowed by default, and sessions can be easily established, but you also can check for User-ID, App-ID, and Content-ID, and perform logging, QoS, decryption, LLDP, zone protection, DoS protection, and NAT.
Security Policy #1
Source Zone
Destination Zone
Logging Profile
Zone 1
Zone 2
VLAN 25 information
(which is a specific customer in this case)
Security Policy #2
Source Zone
Destination Zone
Logging Profile
Zone 3
Zone 4
Host 10.1.2.2 information
(which is a specific server in this case)

To configure a Virtual Wire subinterface, go to **Network > Interfaces > Ethernet** and select, but do not open, a Virtual Wire interface. Then click **Add Subinterfaces** at the bottom of the web interface window:

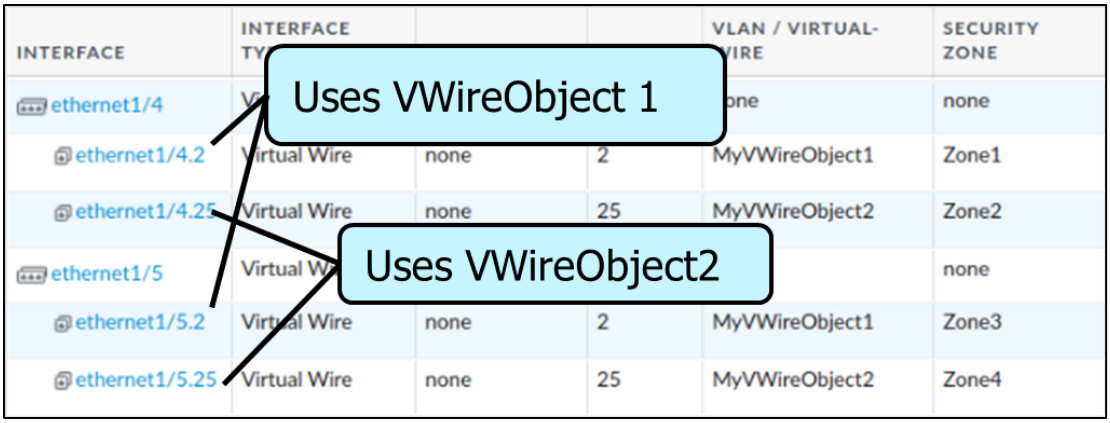

Virtual Wire Subinterface

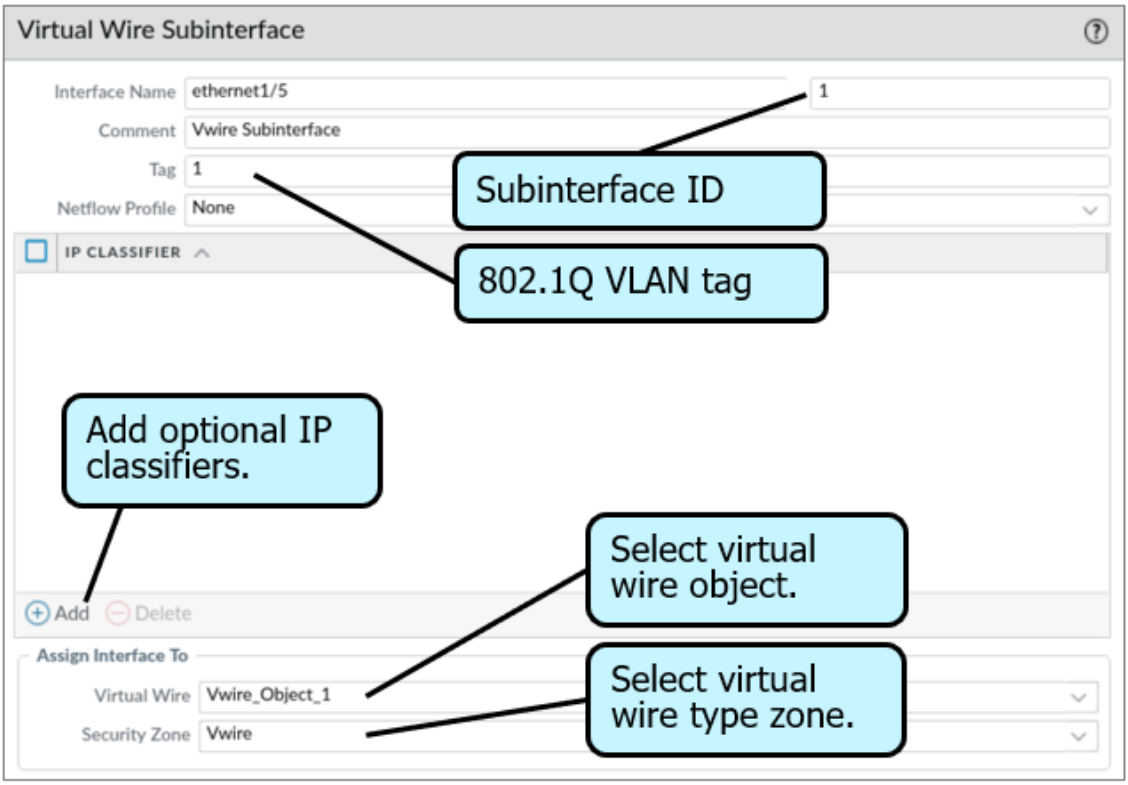

Layer 2 Interfaces

Layer 2 interfaces are used to switch traffic between other Layer 2 interfaces. Before switching can take place, each Layer 2 interface must be assigned to a VLAN object. Assignment of interfaces that belong to the same VLAN but exist in different Layer 2 zones enables you to analyze, shape, manage, and decrypt the traffic. Layer 2 traffic can route to other Layer 3 interfaces using a Layer 3 VLAN interface. Note that Layer 2 interfaces do not participate in spanning tree other than forward BPDUs.

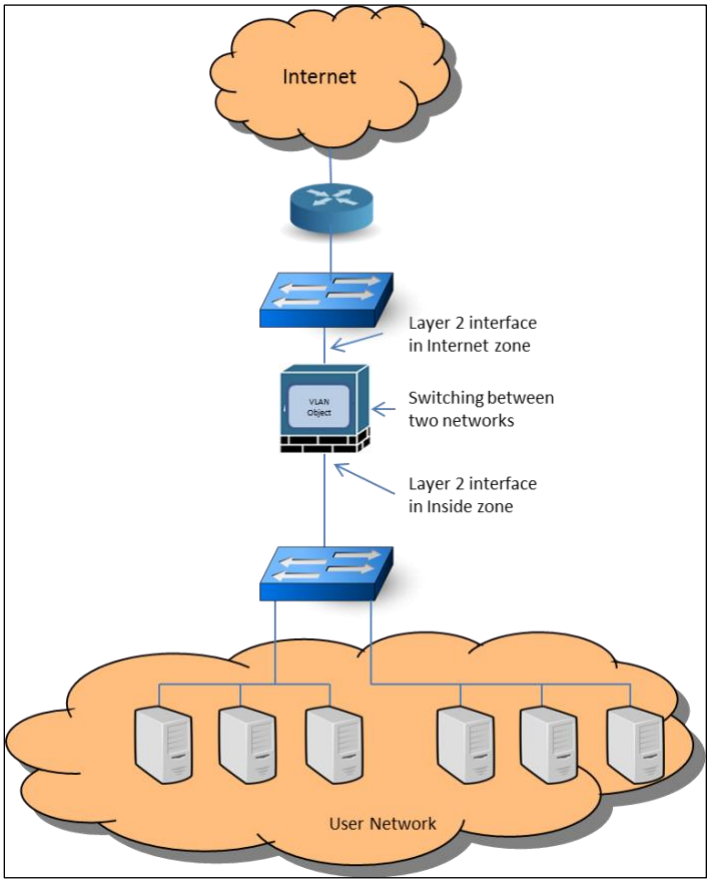

To configure a VLAN object, go to **Network > VLAN > Add**:

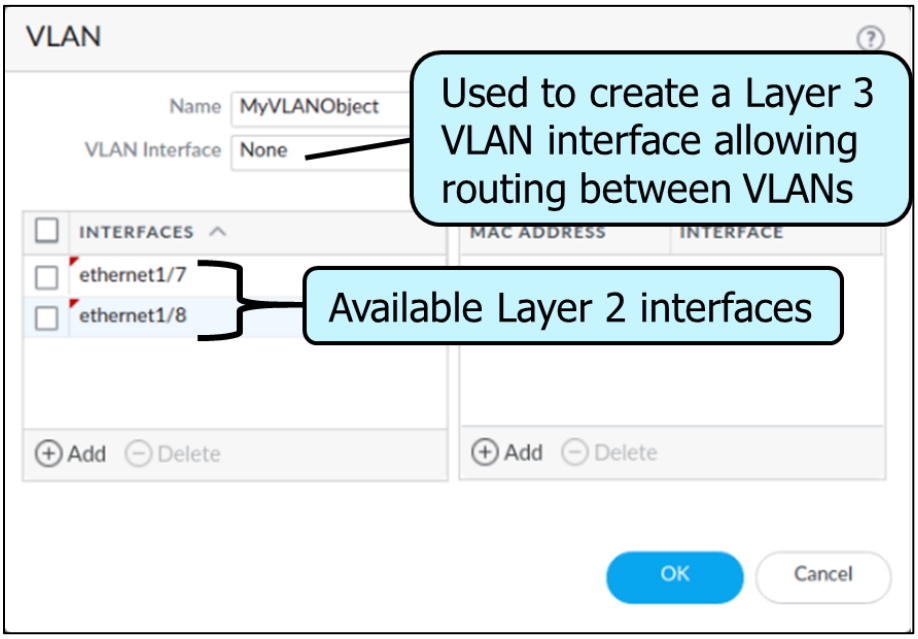

To configure a Layer 2 interface, go to **Network > Interfaces > Ethernet > <select_interface>**:

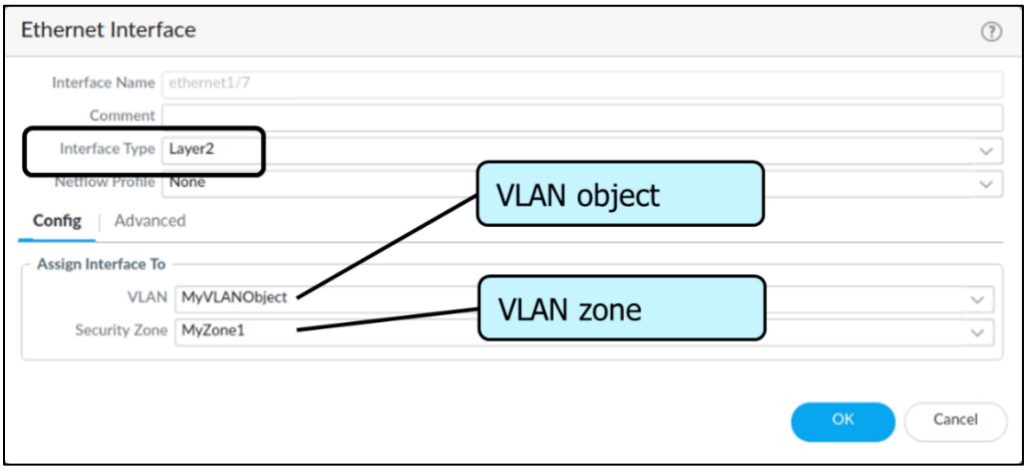

Layer 2 Subinterfaces

Layer 2 interfaces can be subdivided into Layer 2 subinterfaces. For each Ethernet port configured as a physical Layer 2 interface, you can define an additional logical Layer 2 interface (subinterface) for each VLAN tag assigned to the traffic that the port receives. The firewall enables Layer 2 switching between Layer 2 subinterfaces that are connected to the same VLAN object. To enable switching between Layer 2 subinterfaces, assign the same VLAN object to the subinterfaces. Even though Layer 2 subinterfaces are available on a Palo Alto Networks firewall, the best practice is to use Layer 3 subinterfaces. Use of Layer 3 subinterfaces isolates Layer 2 traffic, yet provides routing between subnets.

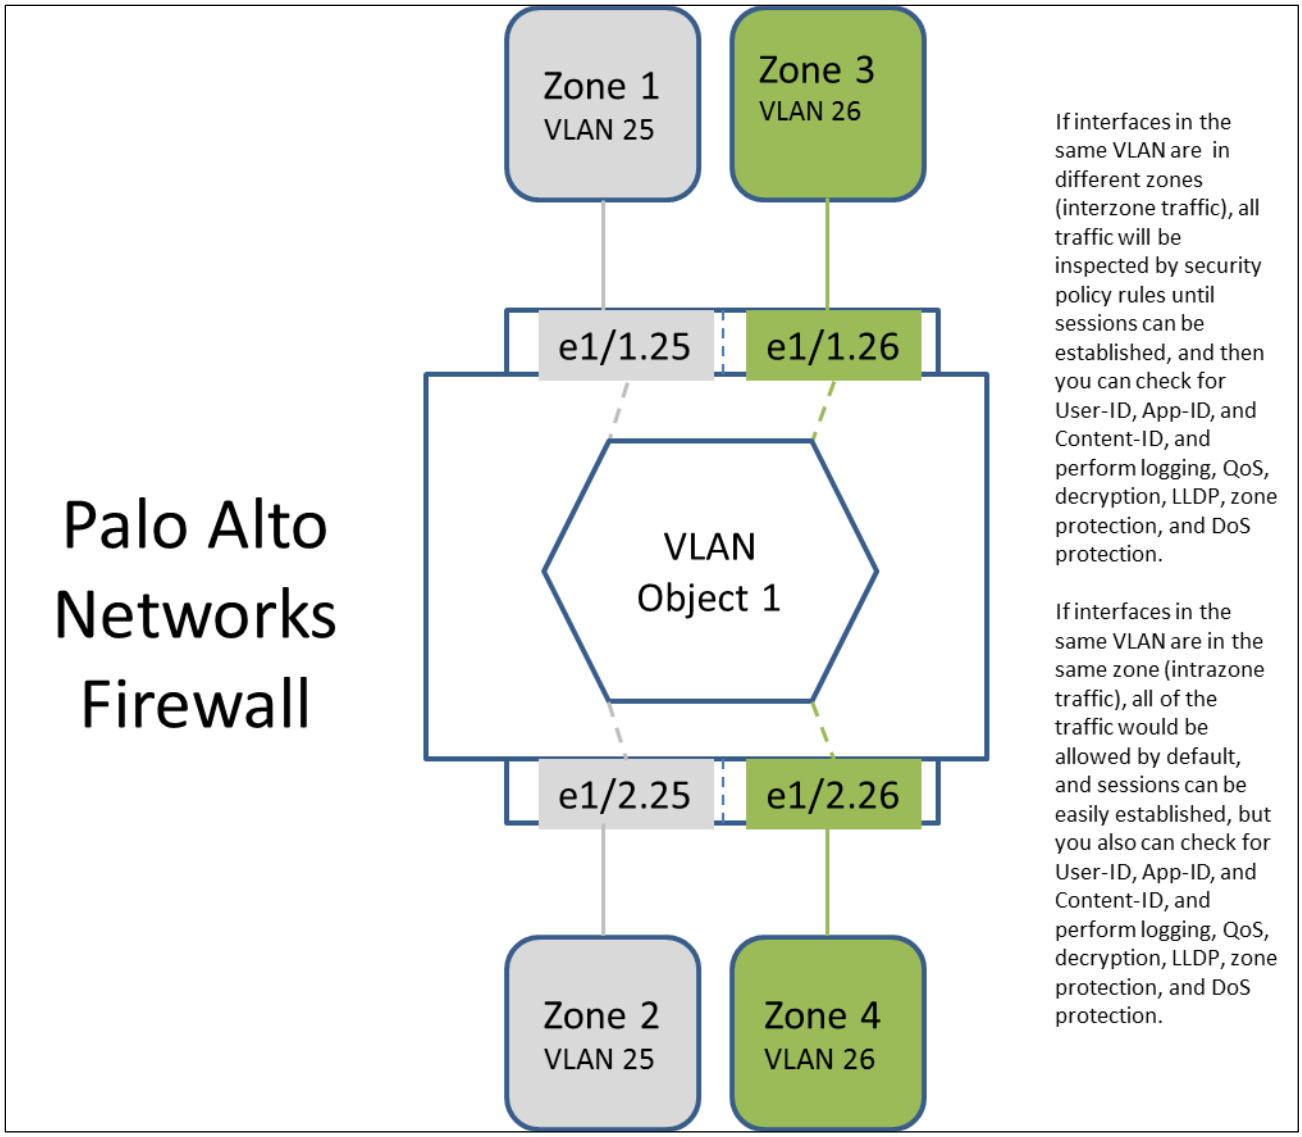

To configure a Layer 2 subinterface, go to **Network > Interfaces > Ethernet > <select_interface>** and select, but do not open, a Layer 2 interface. Then click **Add Subinterfaces** at the bottom of the web interface window:

INTERFACE	INTERFACE TYPE	IP ADDRESS	TAG	VLAN / VIRTUAL-WIRE	SECURITY ZONE	SD-WAN INTERFACE PROFILE
ethernet1/1	Layer2	none	Untagged	MyVLANObject	MyZone1	
ethernet1/1.25	Layer2	none	25	MyVLANObject	MyZone1	
ethernet1/1.26	Layer2	none	26	MyVLANObject	MyZone2	
ethernet1/2	Layer2	none	Untagged	MyVLANObject	none	
ethernet1/2.25	Layer2	none	25	MyVLANObject	MyZone3	
ethernet1/2.26	Layer2	none	26	MyVLANObject	MyZone4	

Layer 2 Subinterface

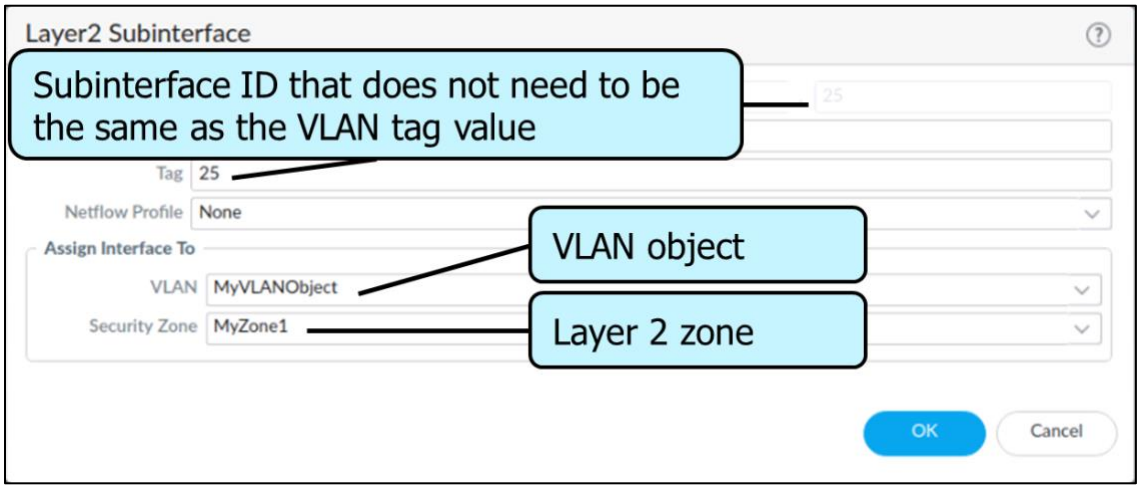

Layer 3 Interfaces

In a Layer 3 deployment, the firewall routes traffic between multiple interfaces. A Virtual Router object must exist for the firewall to route traffic between Layer 3 interfaces. Layer 3 interfaces are assigned IP addresses. PAN-OS software supports both IPv4 and IPv6 addressing. As is the case in most interface types, Layer 3 traffic can be monitored, analyzed, managed, shaped, translated, and encrypted or decrypted. If a tunnel is used for routing or if tunnel monitoring is turned on, the tunnel needs an IP address. The **Advanced** tab contains options that enable you to configure a variety of Layer 3 interface settings such as MTU, static ARP, LLDP, IPv6 NDP, link speed, and duplex settings. Both IPv4 and IPv6 addresses can be configured on a single interface.

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

Unlike Tap, Virtual Wire, or Layer 2 interfaces, Layer 3 interfaces can be used to manage firewalls using an Interface Management Profile. An Interface Management Profile protects the firewall from unauthorized access by defining the protocols, services, and IP addresses that a firewall Layer 3 interface permits for management traffic. Interface Management Profiles are discussed in more detail in a different section of this study guide.

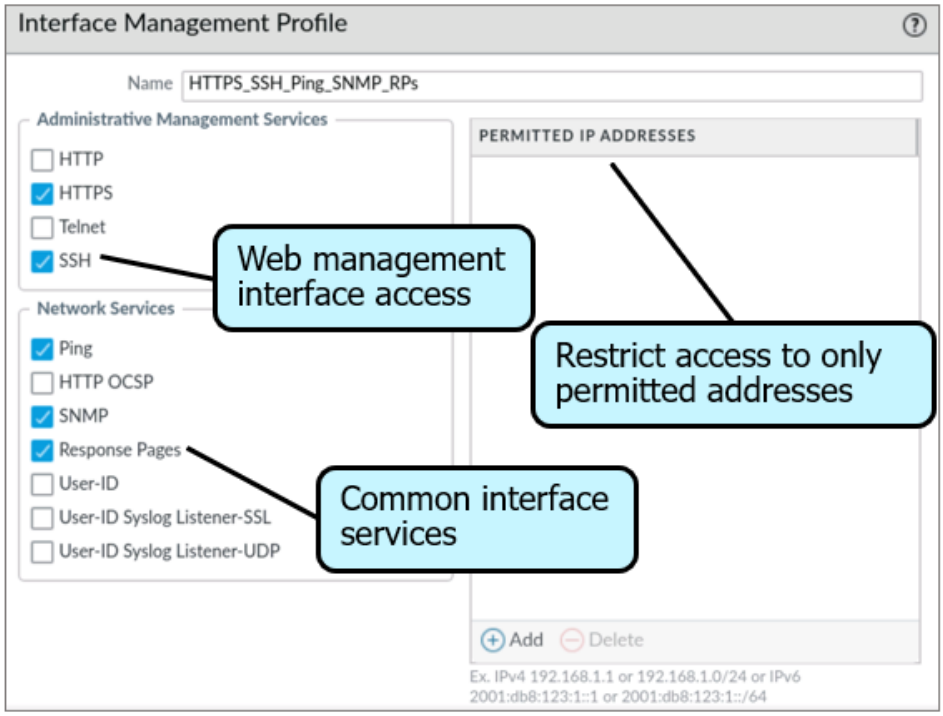

You can configure a Layer 3 interface with one or more static IPv4 addresses or as a DHCP client. A single Layer 3 interface can be assigned multiple IPv4 addresses, although they should not be in the same subnet. You can configure a Layer 3 interface with one or more IPv6 addresses, either as a link-local address or a global address.

Layer 3 interfaces also can be configured as subinterfaces, where each subinterface is assigned a unique IP address.

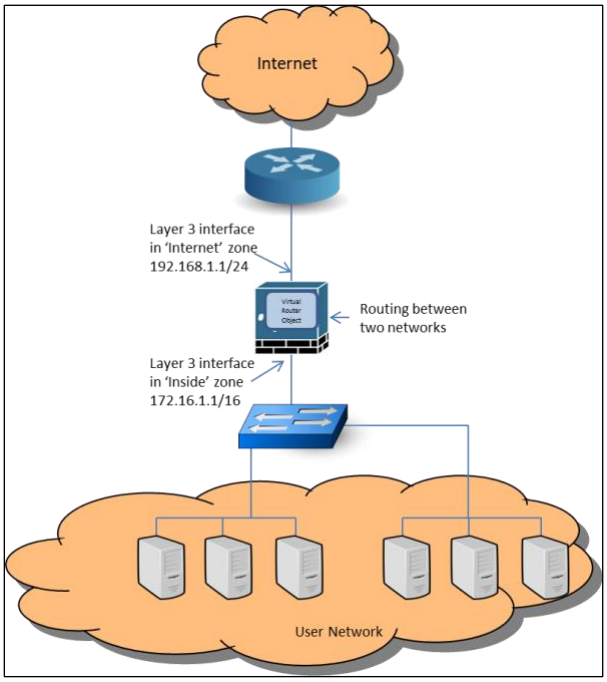

To configure a virtual router object, go to **Network > Virtual Routers > Add**:

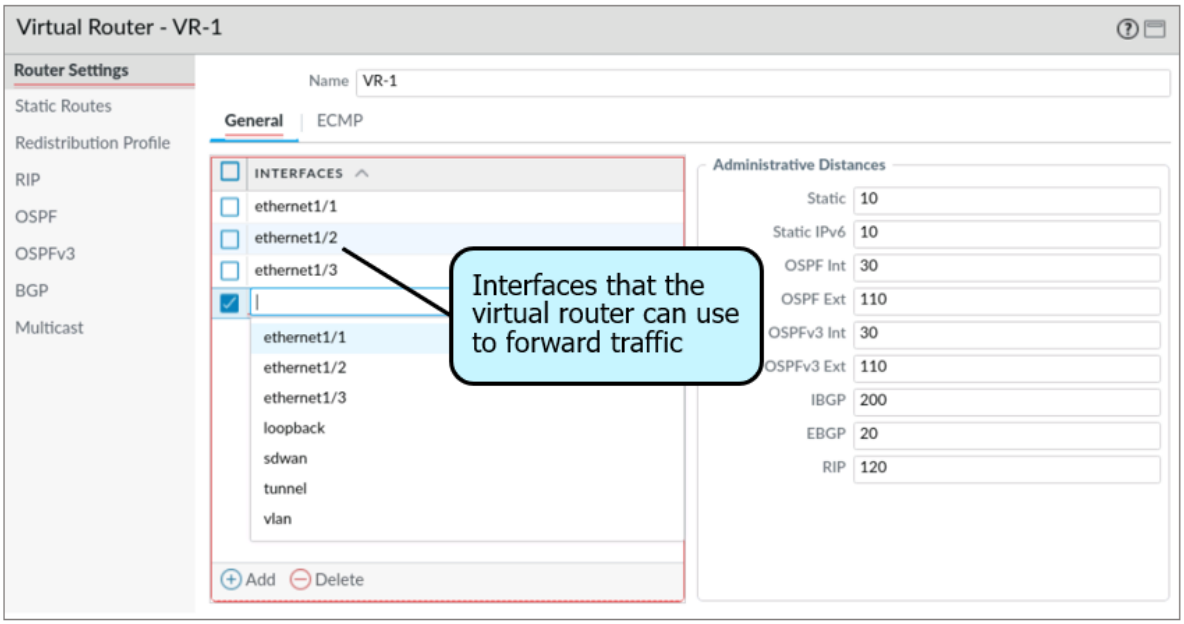

To configure a Layer 3 interface, go to Network > Interfaces > Ethernet > <select_interface>:

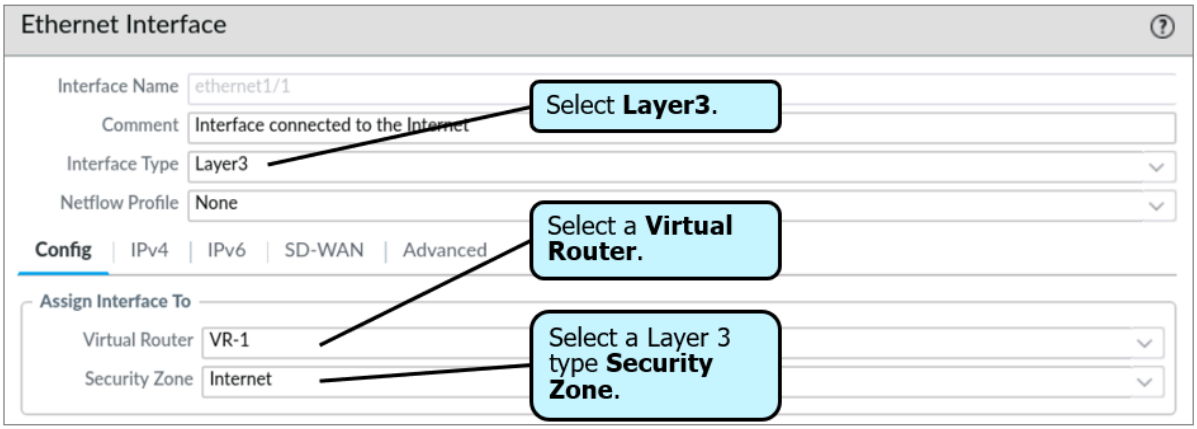

Layer 3 Subinterfaces

For each Ethernet port configured as a physical Layer 3 interface, you can define additional logical Layer 3 interfaces (subinterfaces). Layer 3 subinterfaces possess the same capabilities and features as Layer 3 interfaces, except that Layer 3 subinterfaces are assigned to 802.1Q VLANs. A Virtual Router object is required to route traffic between each VLAN.

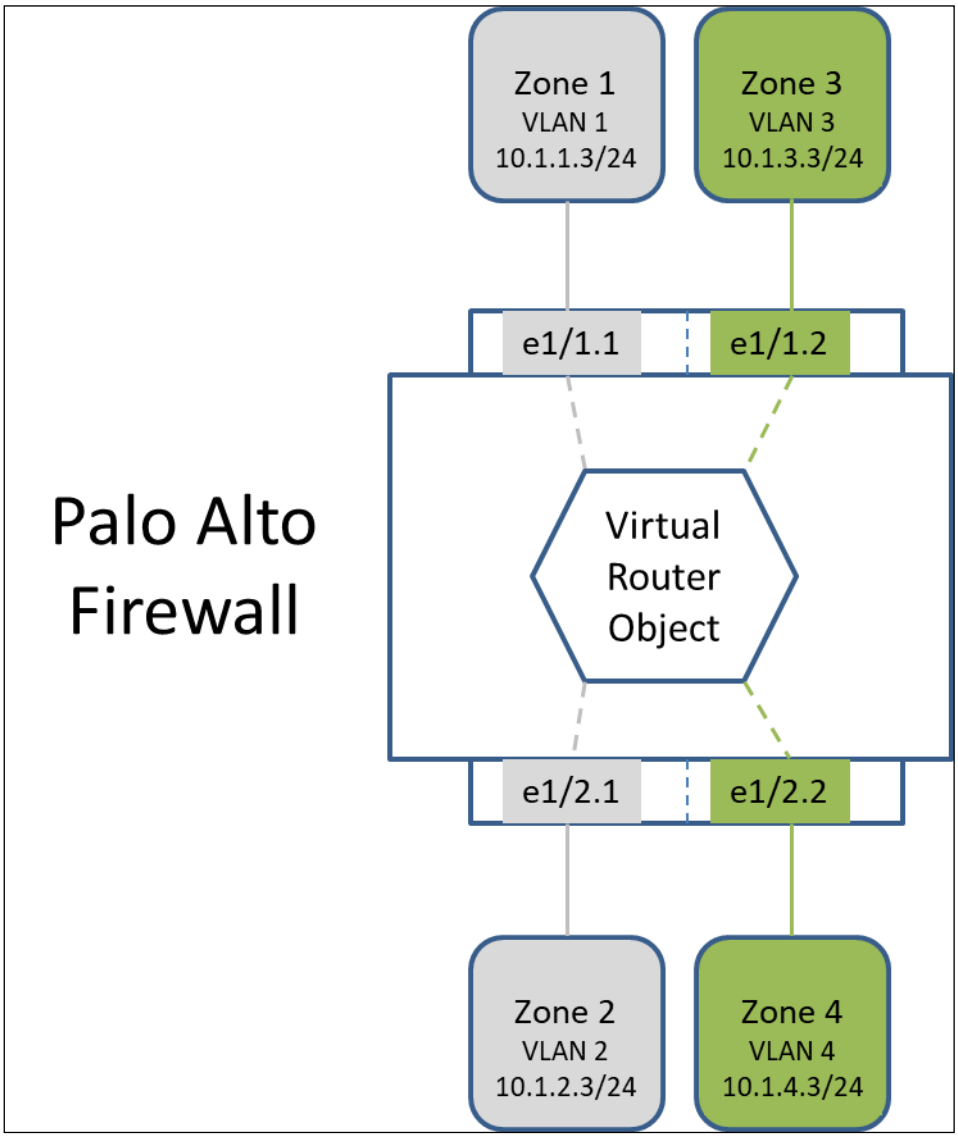

To configure a Layer 3 subinterface, go to **Network > Interfaces > Ethernet**, select a Layer 3 interface, and click **Add Interface**:

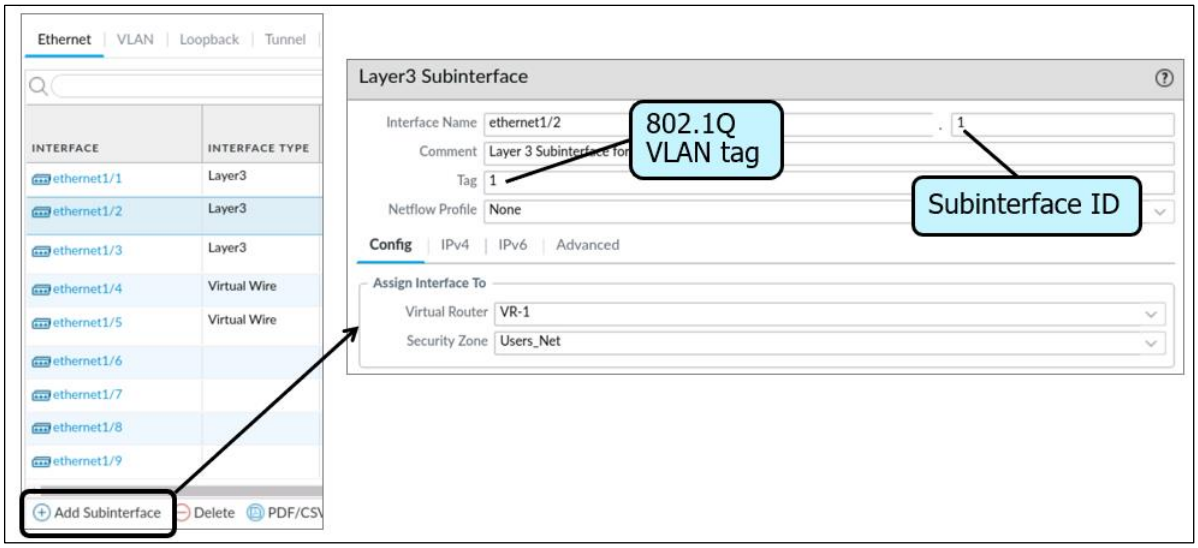

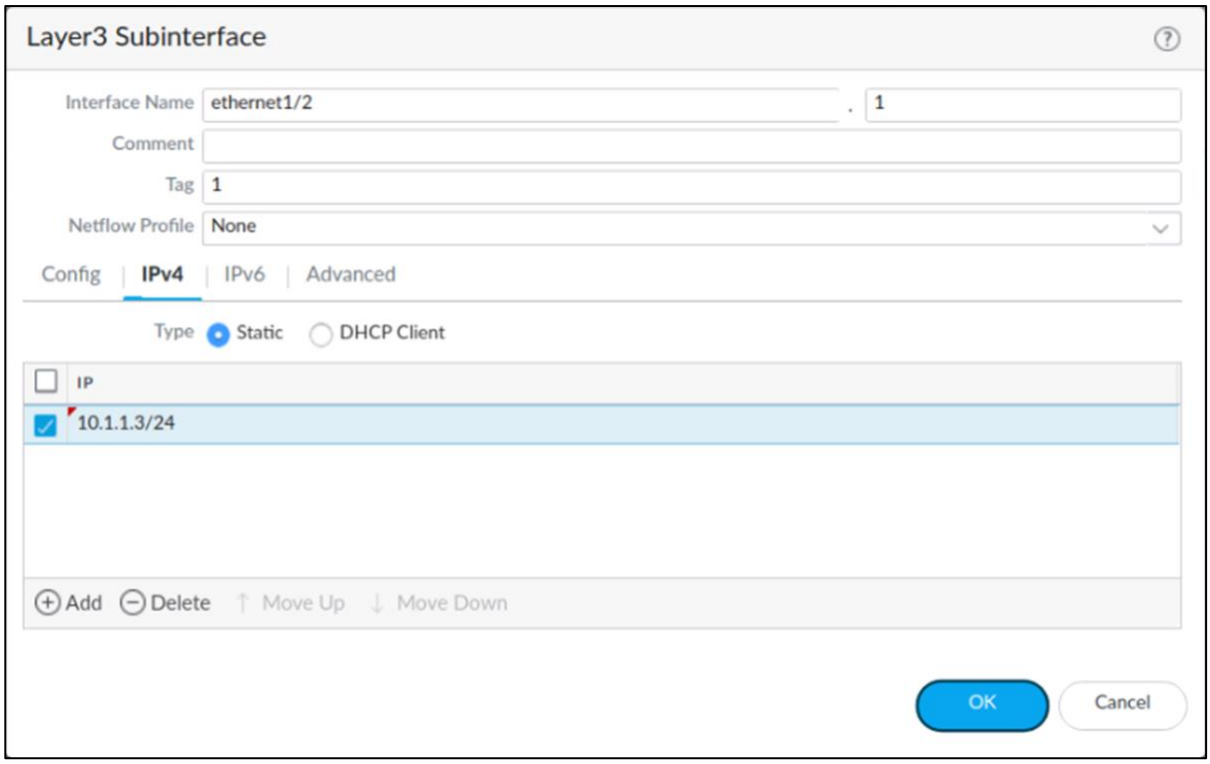

Aggregate Interfaces

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.

Before you configure an aggregate group, you must configure its interfaces. The hardware media can differ among the interfaces assigned to an aggregate group (for example, you can mix fiber optic and copper), but the bandwidth and interface type must be the same

Before you can create an Aggregate Interface, you must first create an Aggregate Interface Group. Select Network > Interfaces > Ethernet and Add Aggregate Group:

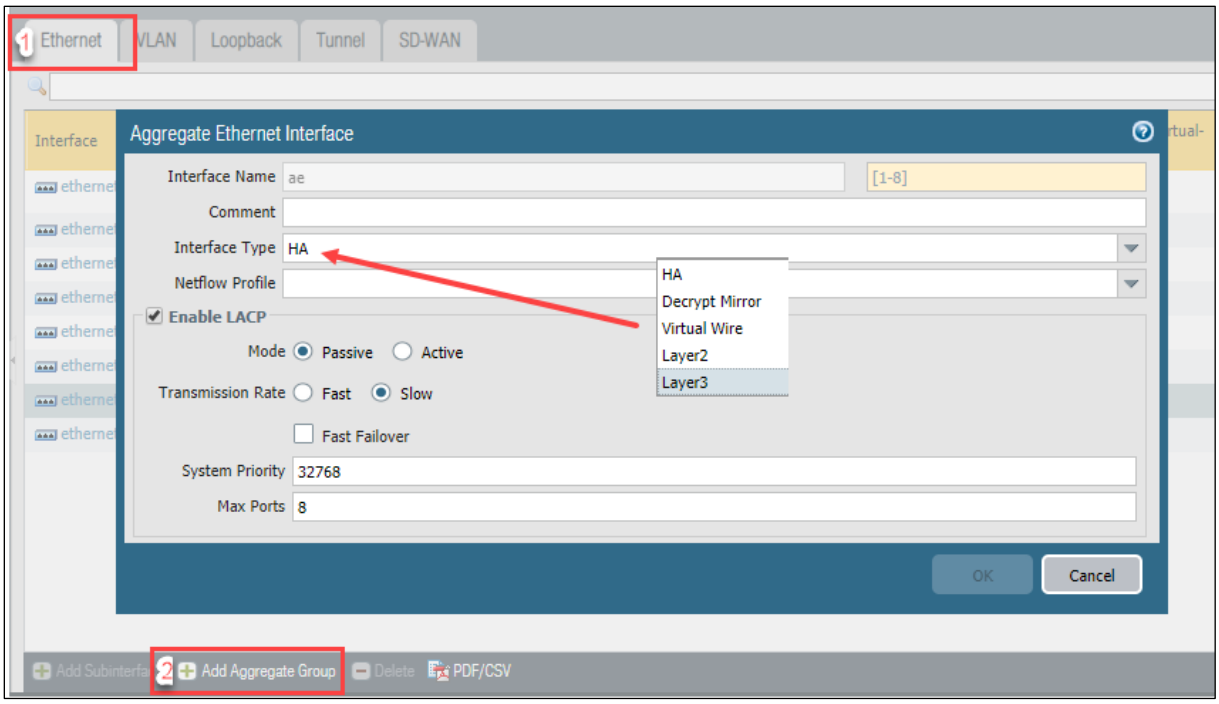

Provide a group ID number and configure LACP protocol settings as required. The **Interface Type** of the Aggregate Interface Group must be chosen from the options shown. After the type is selected, interface configuration might be required. The following screenshot shows the Layer 3 configuration options:

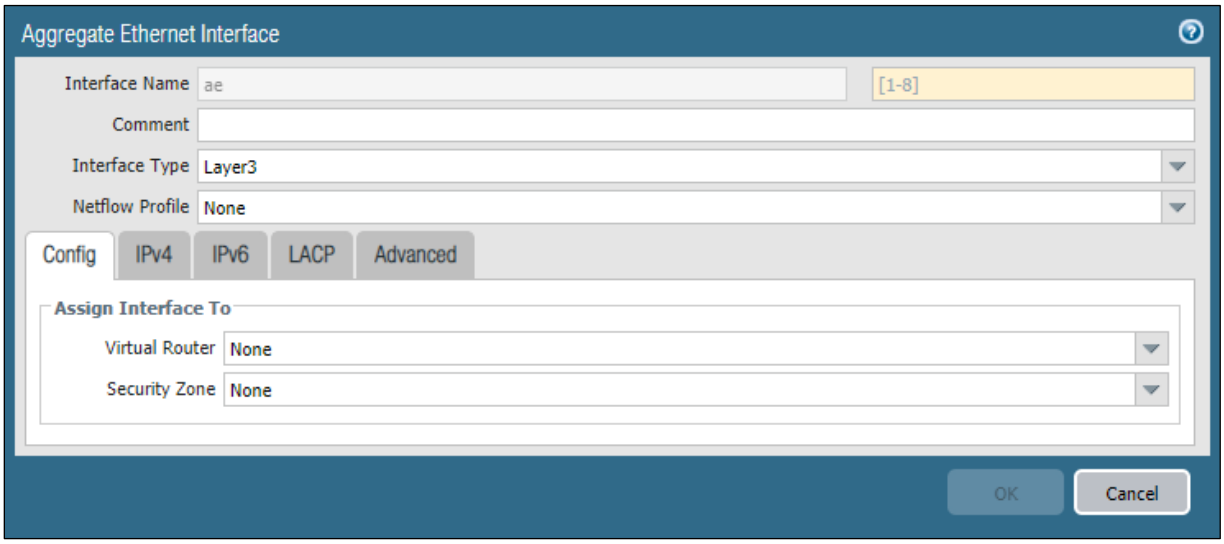

After an Aggregate Interface Group is created, it is added to the available Ethernet interfaces list:

Physical Ethernet interfaces are added to the group by setting the physical **Interface Type** to **Aggregate Ethernet** and selecting the **Aggregate Group** they are assigned to:

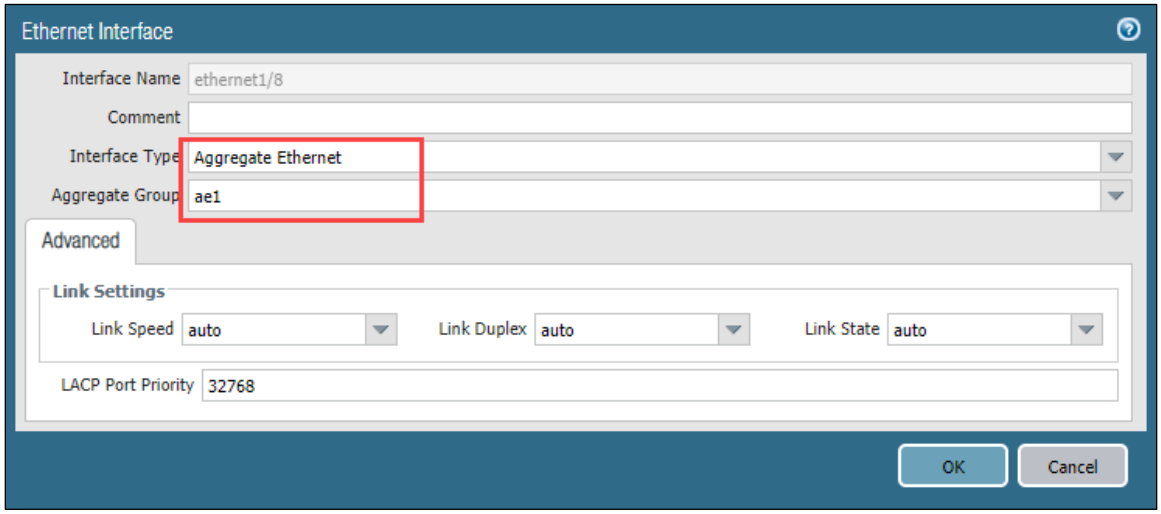

The following screenshot illustrates an Aggregate Interface Group with three assigned physical Ethernet interfaces:

The number of aggregate Ethernet (AE) interface groups that the PA-3200 Series, PA-5200 Series, and most PA-7000 Series firewalls support is 16. All other physical firewall appliances support eight. The exception is the PA-7000 Series firewall with PA-7000-100G-NPC-A and SMC-B, which is 32 AE interface groups. On this firewall, QoS is supported on only the first 16 AE interface groups. On the other supported firewall models, QoS is supported on only the first eight AE interface groups. VM-Series firewalls do not support Aggregate Interface Groups.

Sample Questions

Q1. For inbound inspection, which two actions can be done with a Tap interface? (Choose two.)

a) encrypt traffic

b) decrypt traffic

c) allow or block traffic

d) log traffic

Q2. Which two actions can be done with a Virtual Wire interface? (Choose two.)

a) NAT

b) route

c) switch

d) log traffic

Q3. Which two actions can be done with a Layer 3 interface? (Choose two.)

a) NAT

b) route

c) switch

d) create a Virtual Wire object

Q4. Layer 3 interfaces support which two items? (Choose two.)

a) NAT
b) IPv6
c) switching
d) spanning tree

Q5. Layer 3 interfaces support which three advanced settings? (Choose three.)

a) IPv4 addressing
b) IPv6 addressing
c) NDP configuration
d) link speed configuration
e) link duplex configuration

Q6. Layer 2 interfaces support which three items? (Choose three.)

a) spanning tree blocking
b) traffic examination
c) forwarding of spanning tree BPDUs
d) traffic shaping via QoS
e) firewall management
f) routing

Q7. Which two interface types support subinterfaces? (Choose two.)

a) Virtual Wire
b) Layer 2
c) Loopback
d) Tunnel

Q8. Which two statements are true regarding Layer 3 interfaces? (Choose two.)

a) You can configure a Layer 3 interface with one or more IP addresses as a DHCP client.
b) A Layer 3 interface can only have one DHCP assigned address.
c) You can assign only one IPv4 addresses to the same interface.
d) You can enable an interface to send IPv4 Router Advertisements by selecting the Enable Router Advertisement check box on the Router Advertisement tab.
e) You can apply an Interface Management Profile to the interface.

Q9. Which statement is true regarding aggregate Ethernet interfaces?

a) Members of an Aggregate Interface Group can be of different media types.

b) An Aggregate Interface Group can be set to a type of tap.

c) Member Ethernet interfaces of an Aggregate Interface Group must have the same transmission speeds.

d) A Layer 3 Aggregate Interface Group can have more than one IP assigned to it.

e) Member Ethernet interfaces can be assigned to different virtual routers.

2.8 Configure a virtual router

Virtual Routers

PAN-OS software provides two virtual route engines of which one type can run at a given time: the BGP route engine that supports only BGP and static routing and the legacy route engine that supports multiple dynamic routing protocols. The following firewall models support the BGP route engine:

- PA-7000 Series
- PA-5200 Series
- PA-3200 Series
- VM-Series

Although a supported firewall can have a configuration that uses the legacy route engine and a configuration that uses the BGP route engine, only one route engine is in effect at a time. Each time you change the engine that the firewall will use (enable or disable Advanced Routing to access the BGP route engine or legacy route engine, respectively), you must commit the configuration and reboot the firewall for the change to take effect.

The BGP route engine supports only one logical router (known as a virtual router on the legacy route engine).

Both route engines obtain routes to remote subnets either by the manual addition of static routes or the dynamic addition of routes using dynamic routing protocols. Each Layer 3 Ethernet, Loopback, VLAN, and Tunnel interface defined on the firewall must be associated with a virtual router. Although each interface can belong to only one virtual router, you can configure routing protocols and static routes using either routing engine.

Dynamic routing protocols available on a legacy virtual router are as follows:

- BGP4
- OSPFv2
- OSPVv3
- RIPv2

Multicast routing protocols available on a legacy virtual router are as follows:

- IGMPv1, IGMPv2, IGMPv3
- PIM-SM, PIM-ASM, PIM-SSM

Dynamic routing protocols have administrative distances applied to them that are used to determine the best route to a destination when multiple routes are available from two different routing protocols. The default administrative distances can be modified.

You can create multiple legacy virtual routers, each of which maintains a separate set of routes that aren't shared between these legacy virtual routers, thus enabling you to configure different routing behaviors for different interfaces. Legacy virtual routers can route to other legacy virtual routers within the same firewall if a next hop is specified to reach another legacy virtual router.

The firewall initially populates its learned routes into the firewall's IP routing information base (RIB). The virtual router obtains the best route from the RIB, and then places it in the forwarding information base (FIB). Packets then are forwarded to the next hop router defined in the FIB.

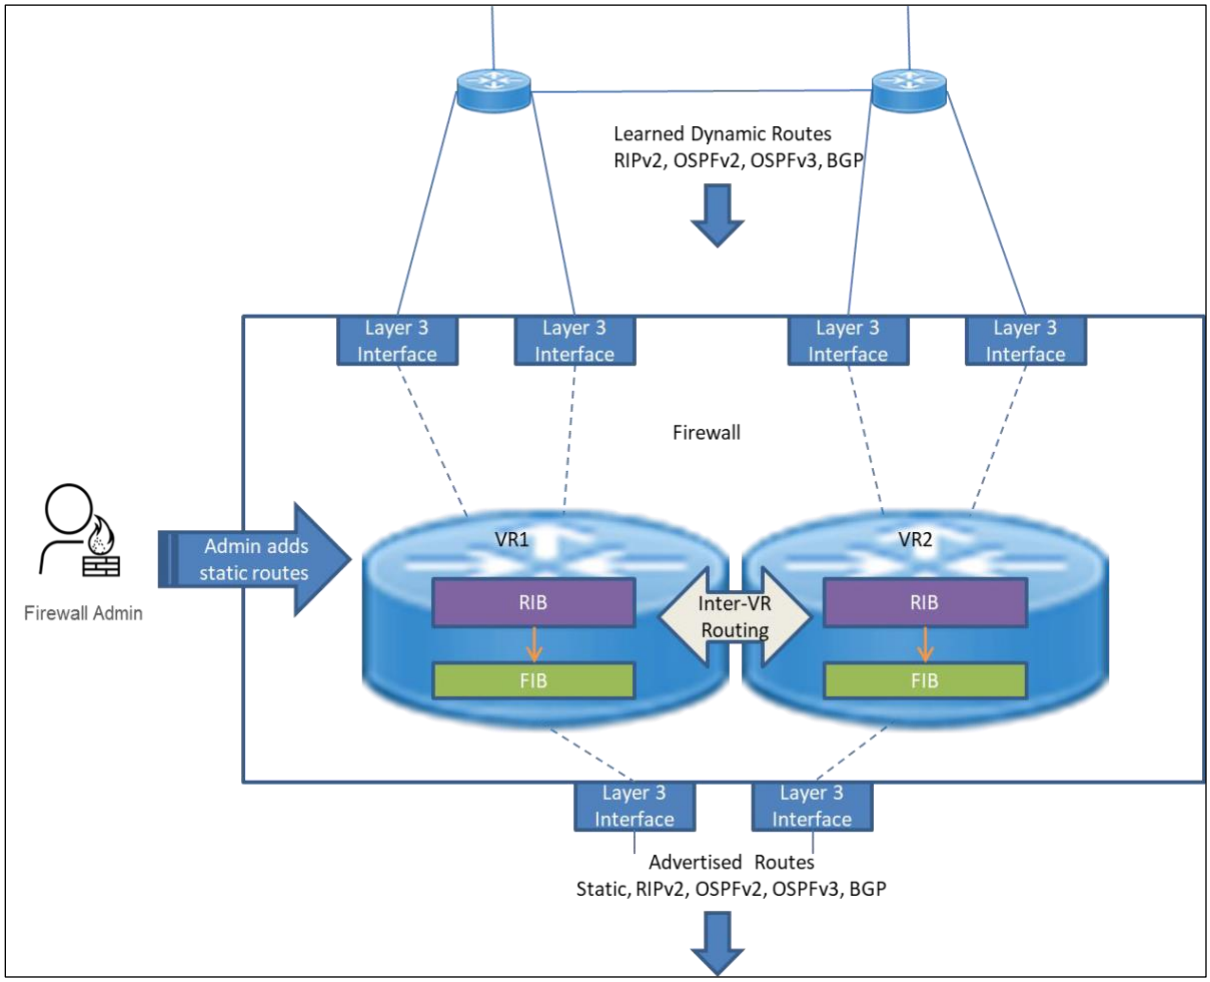

Legacy Virtual Router General Configuration Settings

The administrative distances are shown on the right side of the following screenshot. Most of these distances are consistent with the values in RFCs, but they can be modified to reflect the needs of your environment.

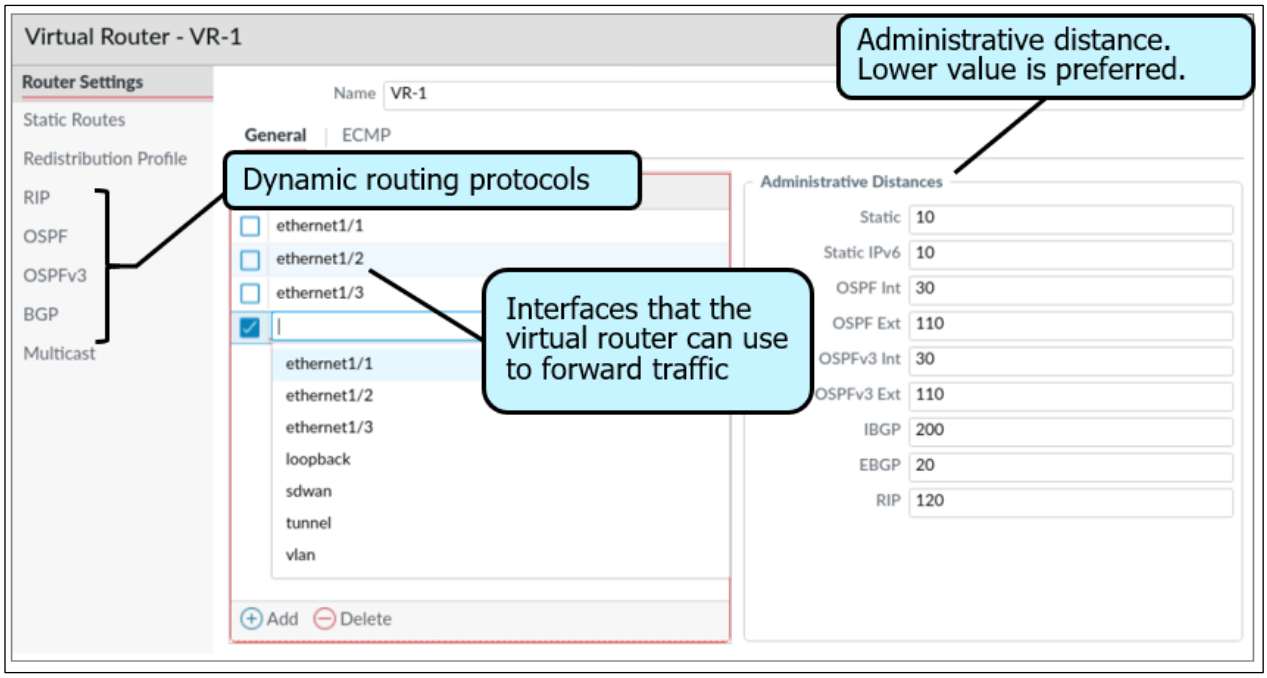

Legacy Static Route Configuration Settings

Static routes have the lowest administrative distances by default, other than locally connected routes. This default administrative distance value is 10, which can be changed.

Static routes have a default metric value of 10, which also can be changed. If you have multiple static routes to the same destination, you can make one preferable over the other by changing the metric. The default metric in the following example was changed from its default value of 10 to 5:

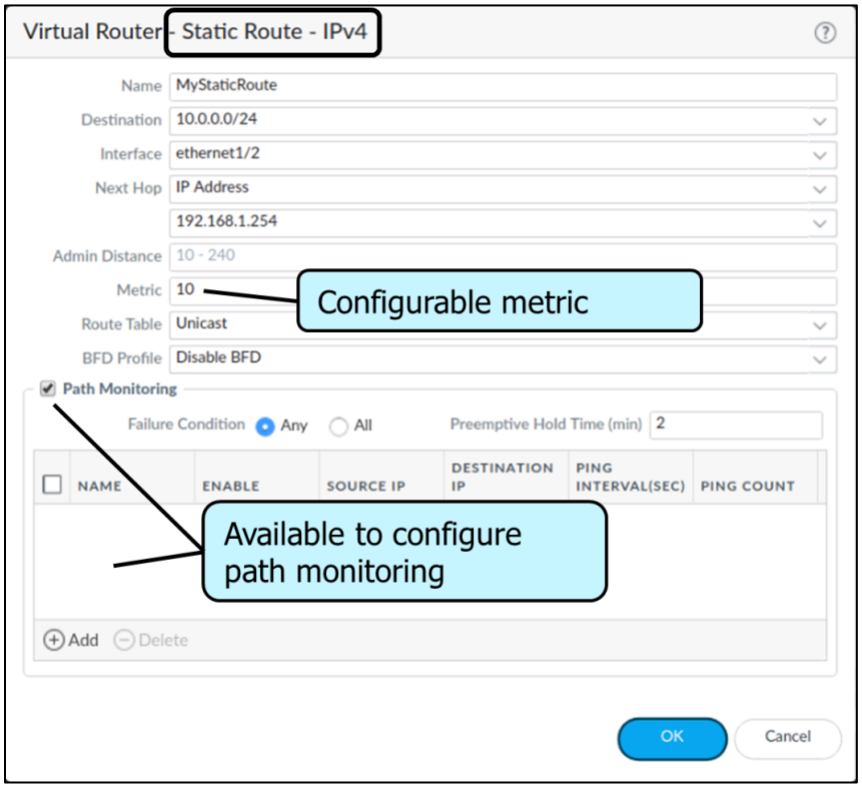

Path Monitoring for Static Routes Configuration Settings

Path monitoring monitors upstream interfaces on remote, reliable devices using ICMP pings. If the path monitoring fails, an associated static route is removed from the routing table. An alternative route then can be used to route traffic.

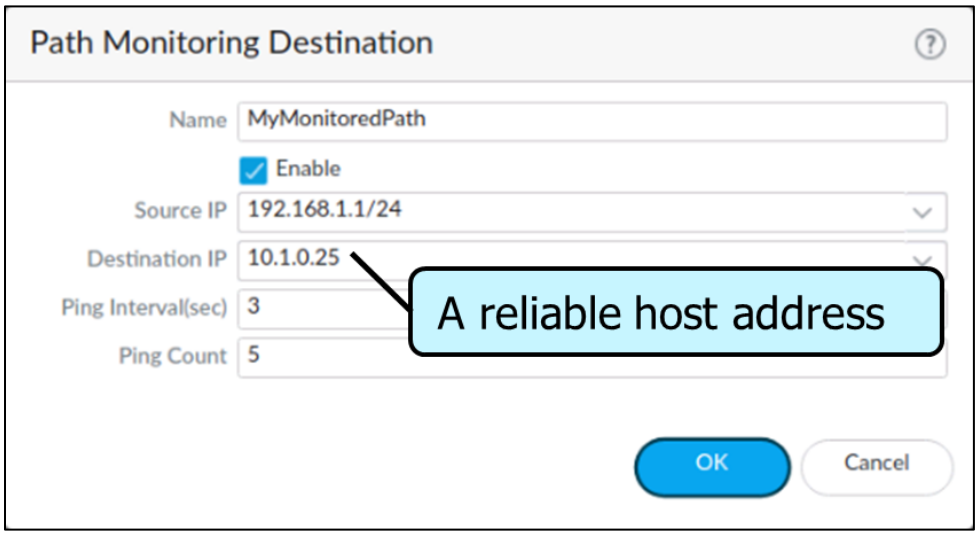

Path Monitoring	Select to enable path monitoring for the static route.
Failure Condition	Select the condition under which the firewall considers the monitored path down and thus the static route down: • **Any**—If any one of the monitored destinations for the static route is unreachable by ICMP, the firewall removes the static route from the RIB and FIB and adds the dynamic or static route that has the next lowest metric going to the same destination to the FIB. • **All**—If all of the monitored destinations for the static route are unreachable by ICMP, the firewall removes the static route from the RIB and FIB and adds the dynamic or static route that has the next lowest metric going to the same destination to the FIB. Select **All** to avoid the possibility of a single monitored destination signaling a static route failure when that monitored destination is simply offline for maintenance, for example.
Preemptive Hold Time (min)	Enter the number of minutes a downed path monitor must remain in Up state—the path monitor evaluates all of its member monitored destinations and must remain Up before the firewall reinstalls the static route into the RIB. If the timer expires without the link going down or flapping, the link is deemed stable, path monitor can remain Up, and the firewall can add the static route back into the RIB. If the link goes down or flaps during the hold time, path monitor fails and the timer restarts when the downed monitor returns to Up state. A **Preemptive Hold Time** of zero causes the firewall to reinstall the static route into the RIB immediately upon the path monitor coming up. Range is 0-1,440; default is 2.
Name	Enter a name for the monitored destination (up to 31 characters).
Enable	Select to enable path monitoring of this specific destination for the static route; the firewall sends ICMP pings to this destination.
Source IP	Select the IP address that the firewall will use as the source 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 appears in the Dynamic IP Interface Status window.
Destination IP	Enter a robust, stable IP address or address object for which the firewall will monitor the path. The monitored destination and the static route destination must use the same address family (IPv4 or IPv6)
Ping Interval (sec)	Specify the ICMP ping interval in seconds to determine how frequently the firewall monitors the path (pings the monitored destination; range is 1-60; default is 3).
Ping Count	Specify the number of consecutive ICMP ping packets that do not return from the monitored destination before the firewall considers the link down. Based on the **Any** or **All** failure condition, if path monitoring is in failed state, the firewall removes the static route from the RIB (range is 3-10; default is 5). For example, a Ping Interval of 3 seconds and Ping Count of 5 missed pings (the firewall receives no ping in the last 15 seconds) means path monitoring detects a link failure. If path monitoring is in failed state and the firewall receives a ping after 15 seconds, the link is deemed up; based on the **Any** or **All** failure condition, path monitoring to **Any** or **All** monitored destinations can be deemed up, and the Preemptive Hold Time starts.

This static route is removed from the routing table until reachability to the next hop is obtained.

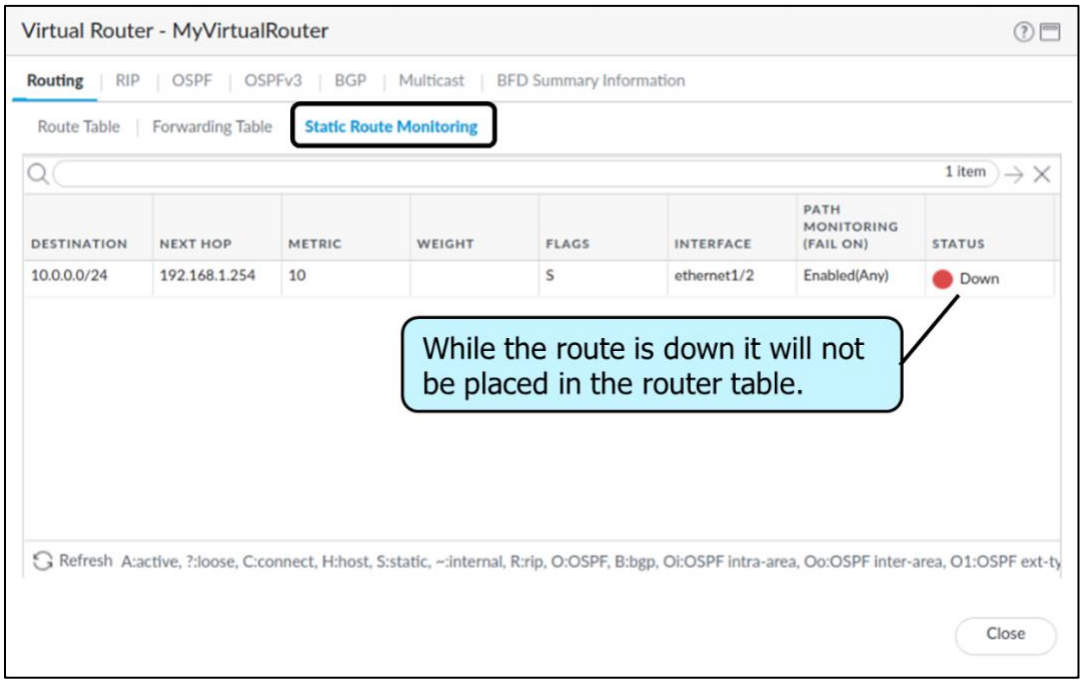

Virtual Router Forwarding Information Base

The following screenshot shows the CLI output of the FIB. A GUI runtime display also is available.

```
admin@firewall-a> show routing fib

total virtual-router shown :              1


--------------------------------------------------------------------------------
virtual-router name: default
interfaces:
   ethernet1/1 ethernet1/2


route table:
flags: u - up, h - host, g - gateway, e - ecmp, * - preferred path

maximum of fib entries for device:                2500
maximum of IPv4 fib entries for device:           2500
maximum of IPv6 fib entries for device:           2500
number of fib entries for device:                 4
maximum of fib entries for this fib:              2500
number of fib entries for this fib:               4
number of fib entries shown:                      4

id      destination           nexthop            flags  interface          mtu
--------------------------------------------------------------------------------
9       10.0.0.0/24           0.0.0.0            u      ethernet1/1        1500
8       10.0.0.1/32           0.0.0.0            uh     ethernet1/1        1500
4       172.16.0.0/24         0.0.0.0            u      ethernet1/2        1500
3       172.16.0.1/32         0.0.0.0            uh     ethernet1/2        1500
--------------------------------------------------------------------------------
```

BGP Route Engine Configuration

To enable the BGP route engine on a firewall, select of **Advanced Routing** at **Device** > **Setup** > **Management** and edit the **General Settings**:

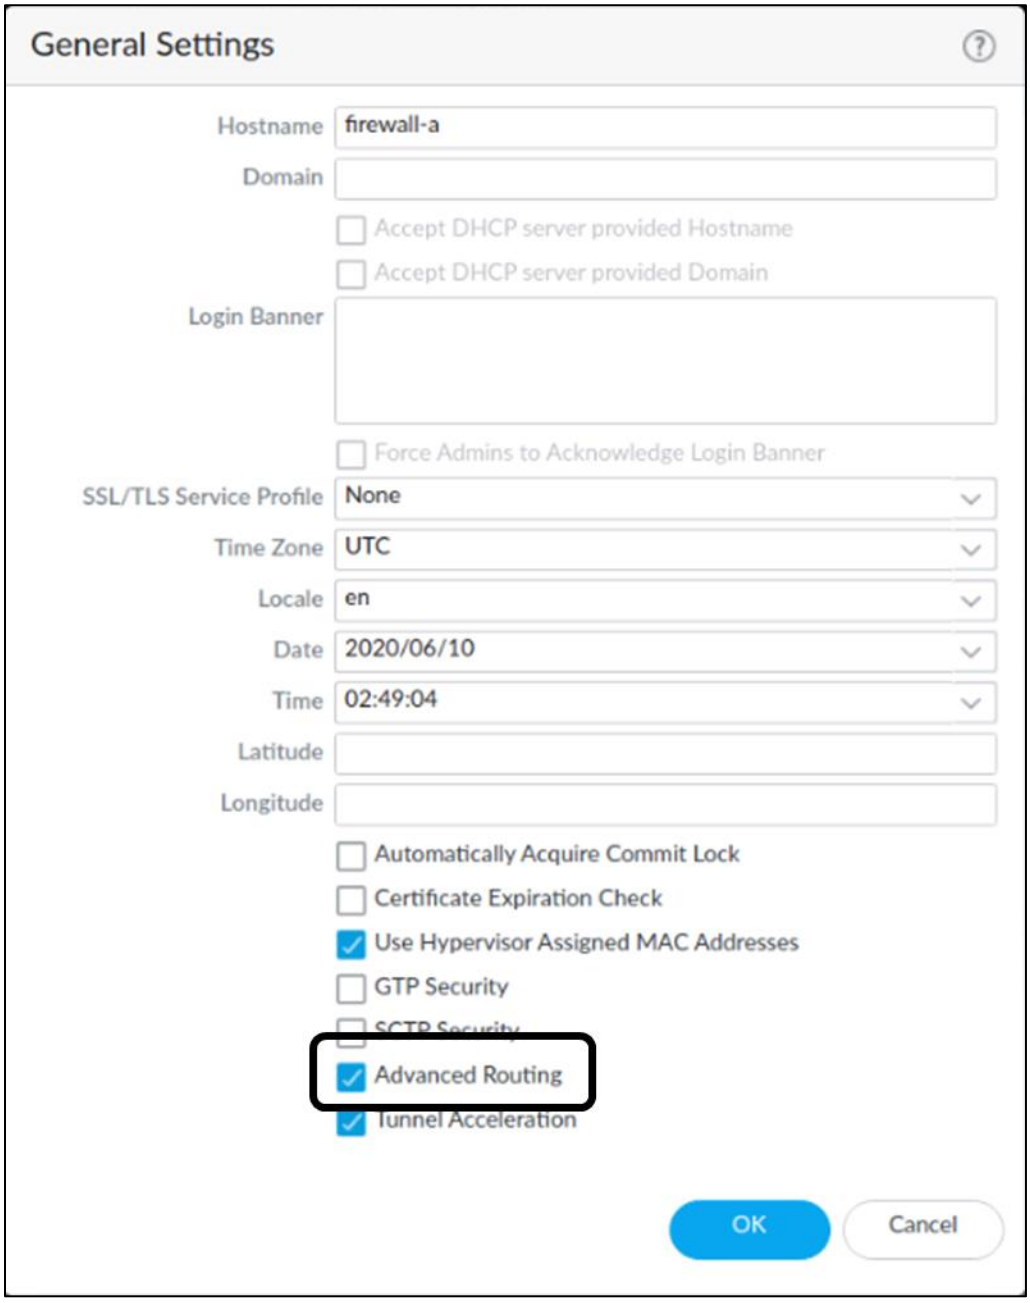

After you make the selection, the firewall must commit the configuration and reboot. After the reboot is complete the BGP routing engine requires the creation of a single **Logical Router** with appropriate settings:

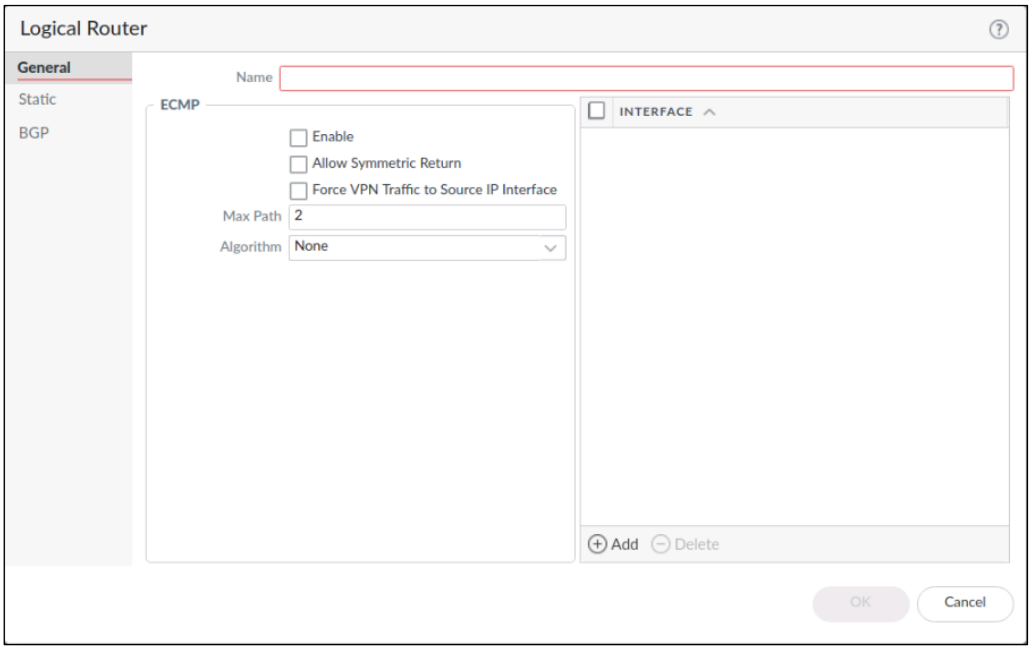

BGP Routing Profiles can be created that specify specific BGP routing feature behavior.

Sample Questions

Q1. What is the default administrative distance of a static route within the PAN-OS software?

a) 1

b) 5

c) 10

d) 100

Q2. Which two dynamic routing protocols are available in the PAN-OS software? (Choose two.)

a) RIP1

b) RIPv2

c) OSPFv3

d) EIGRP

Q3. Which value is used to distinguish the preference of routing protocols?

a) metric
b) weight
c) distance
d) cost
e) administrative distance

Q4. Which value is used to distinguish the best route within the same routing protocol?

a) metric
b) weight
c) distance
d) cost
e) administrative distance

Q5. In path monitoring, what is used to monitor remote network devices?

a) ping
b) SSL
c) HTTP
d) HTTPS
e) link state

Domain 3 – Managing Objects

3.1 Identify how to create address objects

There are four types of address objects:

- IP Netmask
- IP Range
- IP Wildcard Mask
- FQDN

Both IPv4 or IPv6 addresses are supported for the IP Netmask, IP Range, or FQDN address object types. However, IP Wildcard Mask can specify only IPv4 addresses.

To create an address object, perform the following steps:

1. Select **Objects** > **Addresses** and **Add** an address object by **Name**. The name is case-sensitive, and the name must be unique. There is a limit of 63 characters (letters, numbers, spaces, hyphens, and underscores).
2. Select the **Type** of address object.
3. Enter a **tag** to apply to the address object.
4. **Commit** changes.
5. View logs filtered by your address object.
6. View a custom report based on your address object.
7. Use a filter in the ACC to view network activity. Select **ACC > Network Activity**.

References

- Use Tags to Group and Visually Distinguish Objects https://docs.paloaltonetworks.com/pan-os/10-0/pan-os-admin/policy/use-tags-to-group-and-visually-distinguish-objects.html
- Register IP Addresses and Tags Dynamically https://docs.paloaltonetworks.com/pan-os/10-0/pan-os-admin/policy/register-ip-addresses-and-tags-dynamically.html

3.2 Identify how to create services

Administrator accounts control access to firewalls. By default, Palo Alto Networks firewalls have a predefined, default local admin account that has full access. Administrator accounts can be either local (internal) or non-local (external). Additional local or external administrator accounts can be created with customized administrative privileges by assigning them to **Role Based** admin role profiles, or you can assign administrator accounts to built-in account types using **Dynamic** admin roles.

Administrative Role Types

A *role* defines the type of access that an administrator has to the firewall. The two role types are **Role Based** profile roles and **Dynamic** roles:

- **Role Based** profile roles: These are 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 role 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. After 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.

Administrator Account Configuration

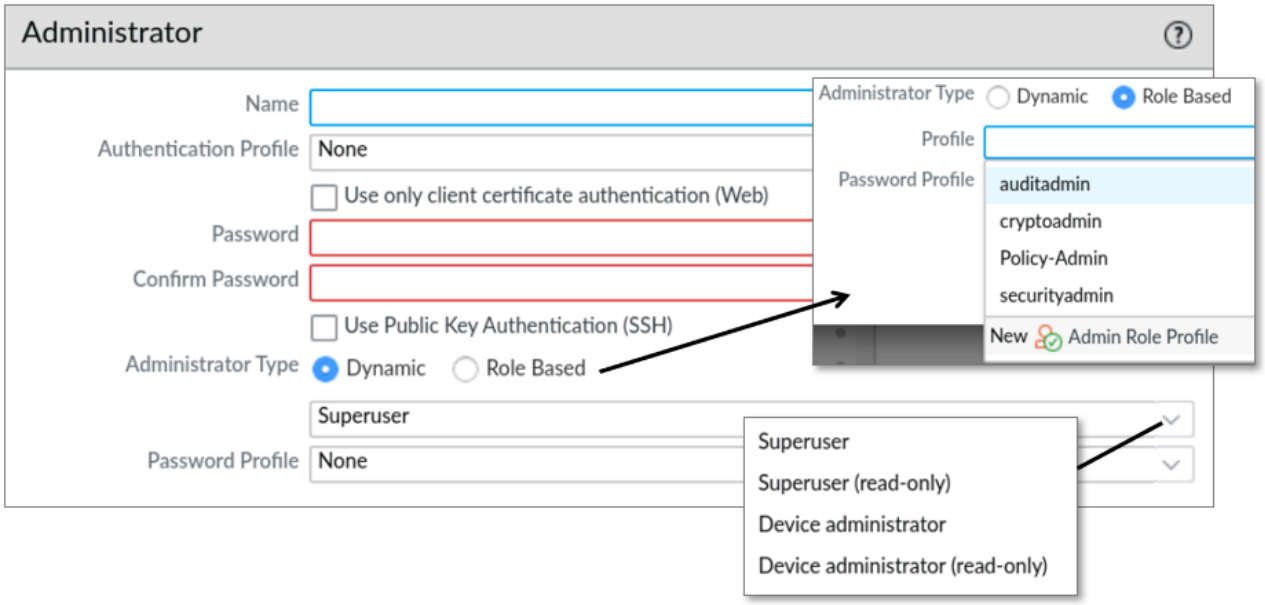

Role-Based Profile Types

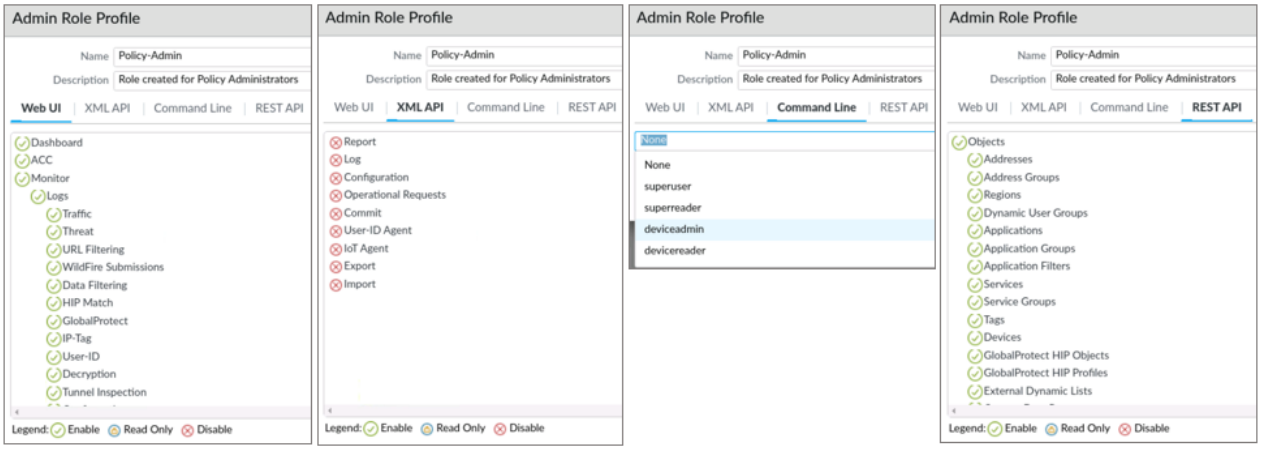

- **Dynamic** roles: These are built-in or predefined roles that provide access to the firewall. When new features are added, the firewall automatically updates the definitions of **Dynamic** roles; you never need to manually update them. The following list identifies the access privileges associated with dynamic roles:
    - **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, available on only firewalls that support virtual systems
    - **Virtual system administrator (read-only**): Read-only access to a selected vsys on the firewall, available on only firewalls that support virtual systems
    - **Device administrator**: Full access to all firewall settings except for defining new accounts or virtual systems
    - **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)

Administrator Account Configuration

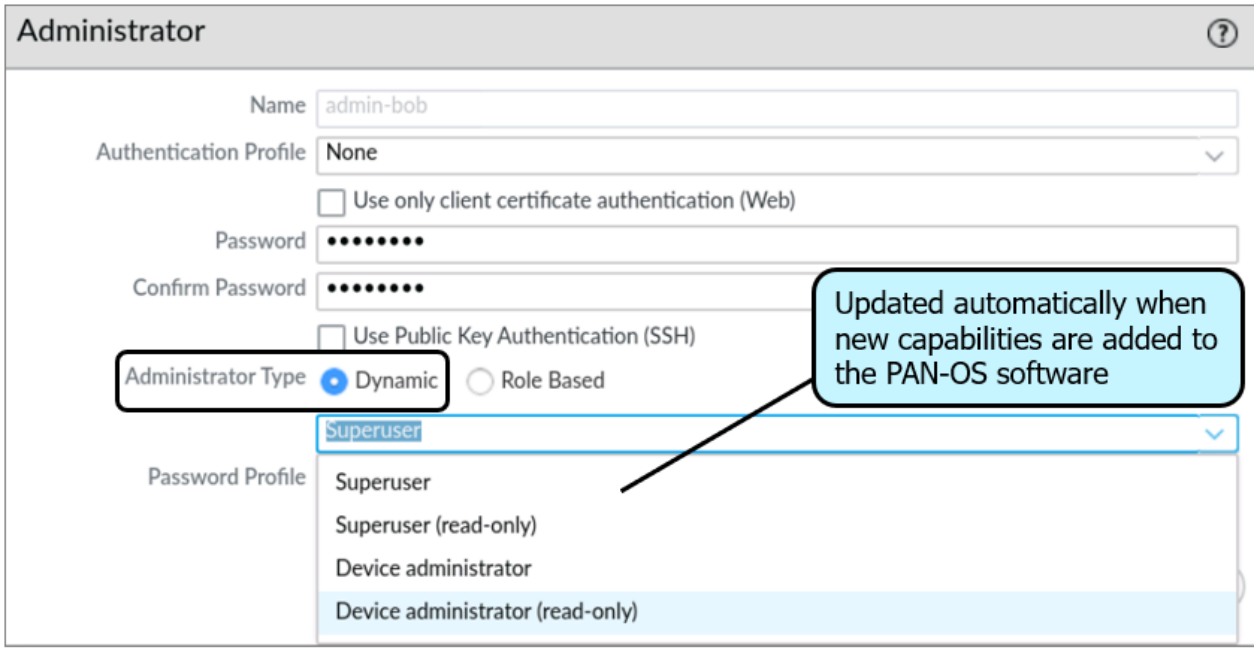

Local administrator accounts are authenticated using a local database:

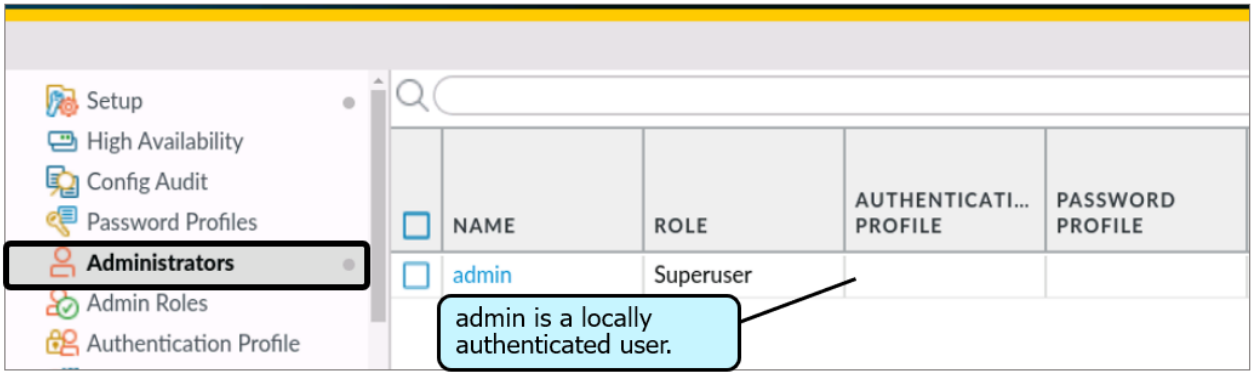

External administrator accounts require an external authentication service that is specified using an Authentication Profile.

PAN-OS software supports the following authentication types:

- None
- Local Database
- RADIUS
- LDAP
- TACACS+

- SAML
- Kerberos

Authentication Profiles provide authentication settings that you can apply to administrator accounts, SSL-VPN access, and Captive Portal. An Authentication Profile configuration screenshot follows:

Authentication Profiles

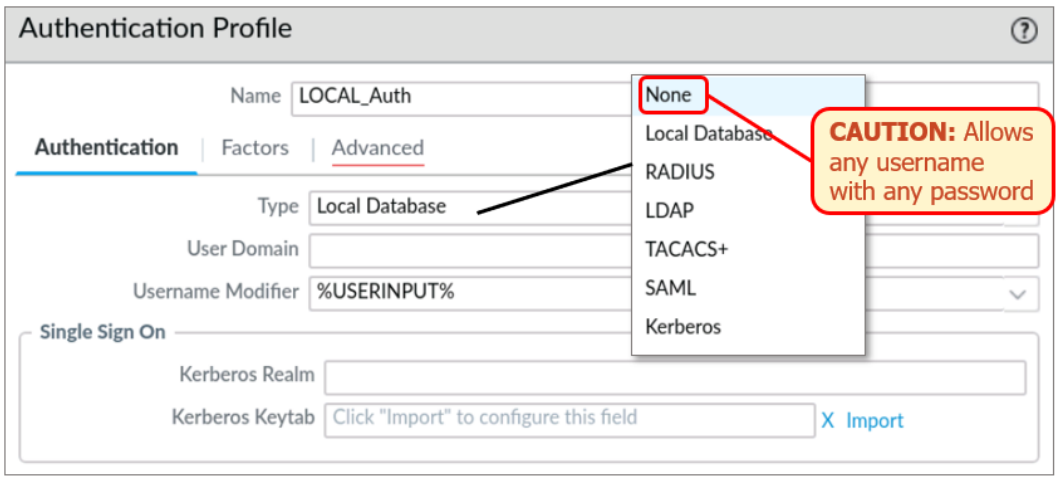

An Authentication Profile references a Server Profile:

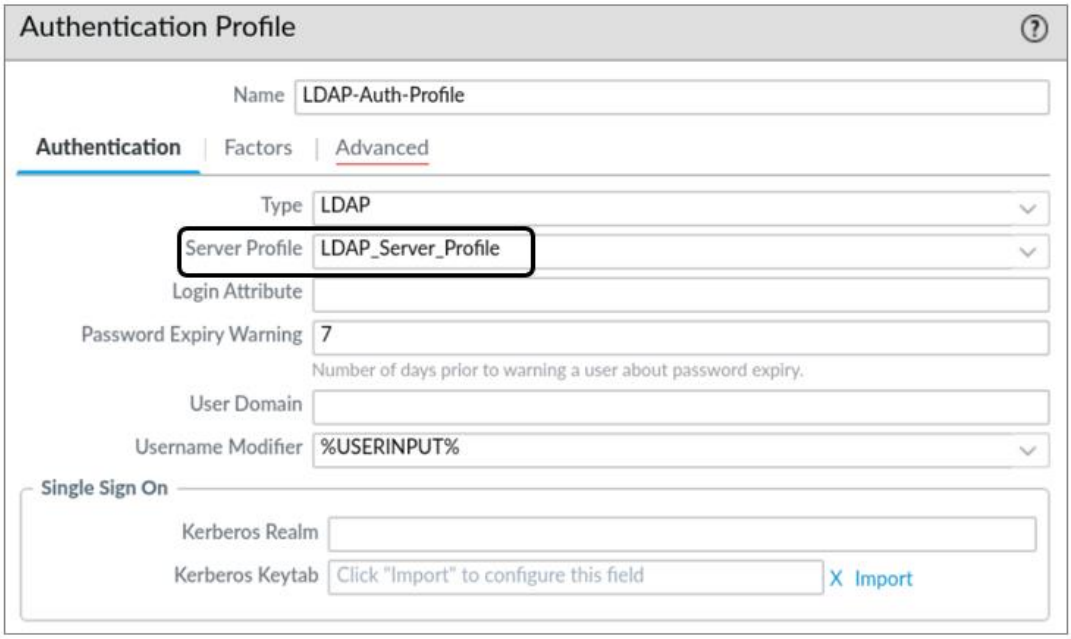

A Server Profile includes the server name, its IP address, the service port that it is listening to, and other values. An example of a RADIUS Server Profile follows:

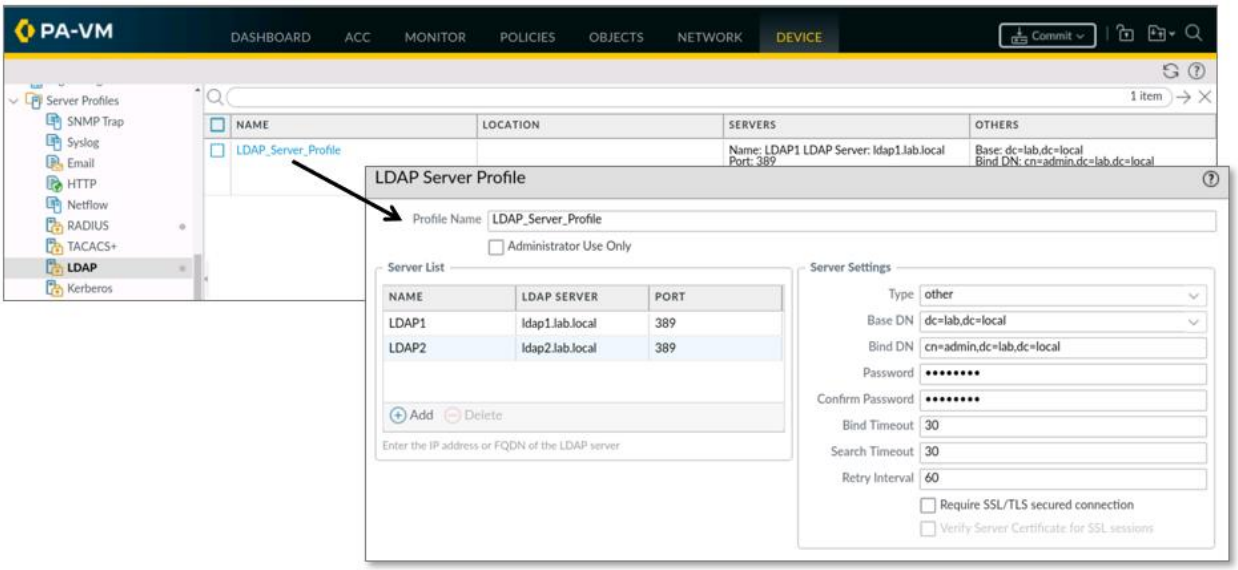

Authentication Sequence

Admin Roles for external administrator accounts can be assigned to an Authentication Sequence, which includes a sequence of one or more Authentication Profiles that are processed in a specific order. The firewall checks against each Authentication Profile within the Authentication Sequence until one Authentication Profile successfully authenticates the user. If an external administrator account does not reference an Authentication Sequence, it directly references an Authentication Profile instead. A user is denied access only if authentication fails for all the profiles in the Authentication Sequence. A depiction of an Authentication Sequence follows:

Authentication Sequence

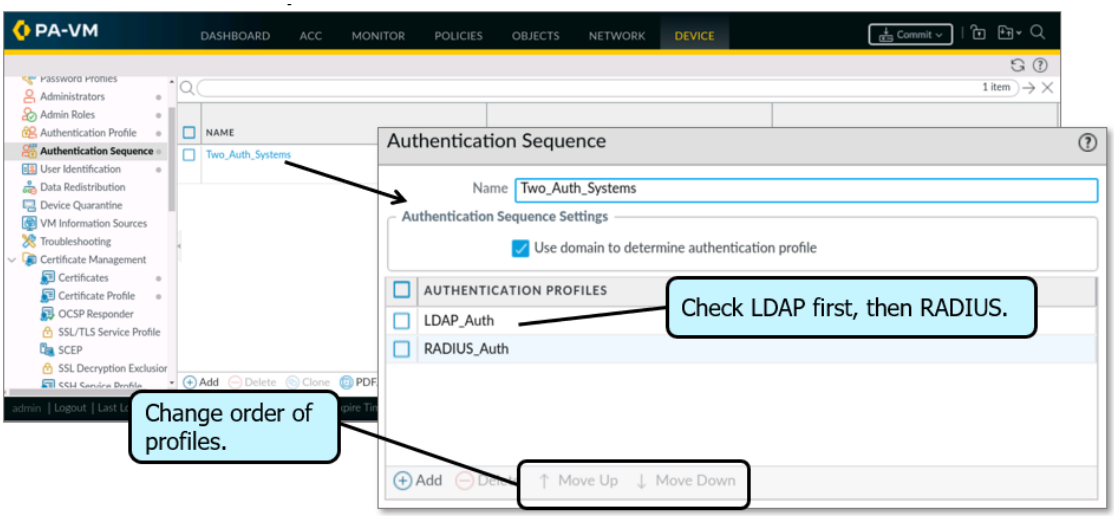

Administrator Account Passwords

To ensure tighter security, you should enable Minimum Password Complexity Requirements. These global settings are applied to all local administrator accounts and help protect the firewall against unauthorized access for administrator accounts that require stricter complexity and aging requirements than do accounts for standard administrators.

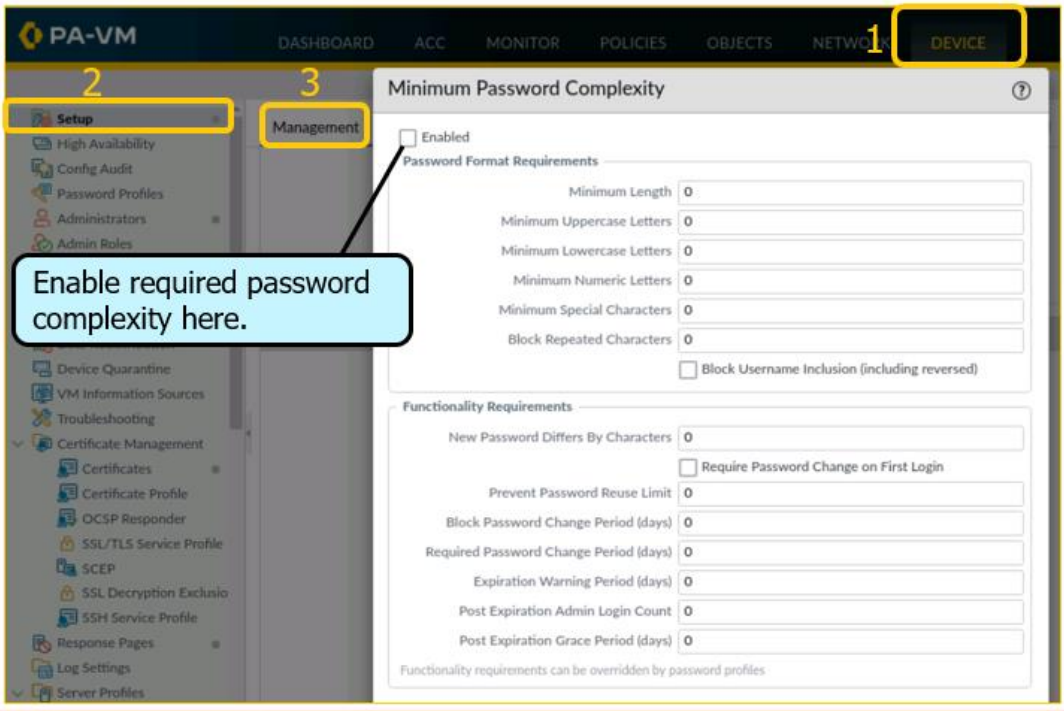

A password profile can be assigned to a local administrator account, which overrides the global password settings:

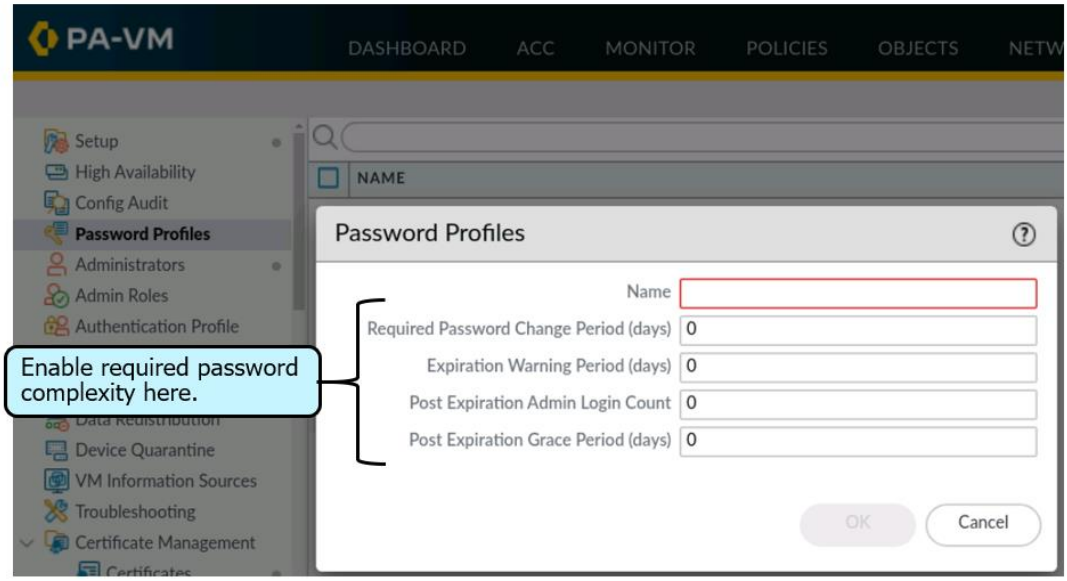

Configuration Logs

Configuration 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.

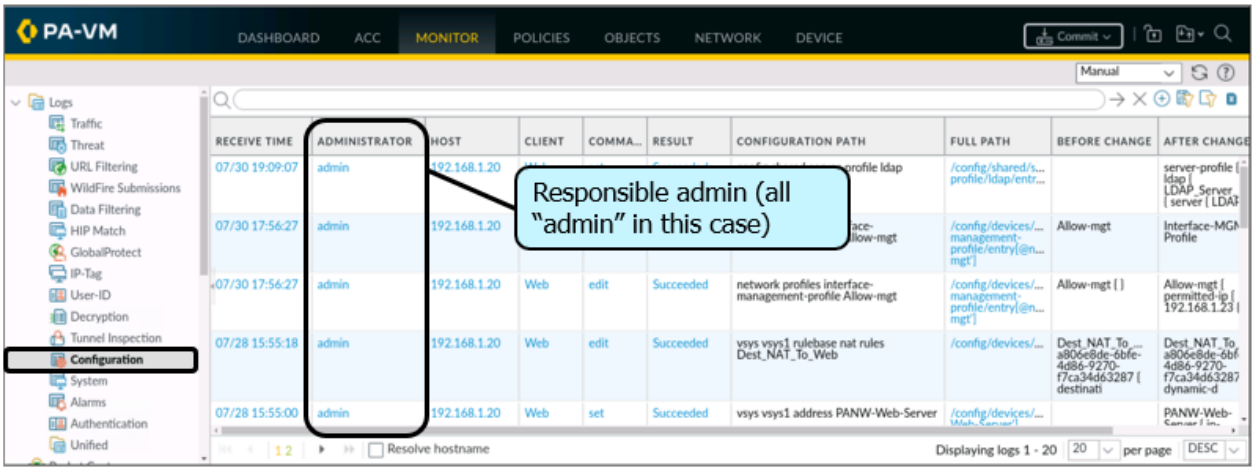

Sample Questions

Q1. Which two statements are true about a Role Based Admin Role Profile role? (Choose two.)

a) It is a built-in role.

b) It can be used for CLI commands.

c) It can be used for XML API.

d) Superuser is an example.

Q2. The management console supports which two authentication types? (Choose two.)

a) RADIUS

b) SMB

c) LDAP

d) TACACS+

e) AWS

Q3. Which two Dynamic Admin Role types are available on the PAN-OS software? (Choose two.)

a) superuser
b) superuser (write only)
c) device user
d) device administrator (read-only)

Q4. Which type of profile does an Authentication Sequence include?

a) Security
b) Authorization
c) Admin
d) Authentication

Q5. An Authentication Profile includes which other type of profile?

a) Server
b) Admin
c) Customized
d) Built-In

Q6. True or false: Dynamic Admin Roles are called “dynamic” because you can customize them.

a) true
b) false

Q7. Which profile is used to override global Minimum Password Complexity Requirements?

a) Authentication
b) Local
c) User
d) Password

3.3 Identify how to use predefined Palo Alto Networks external dynamic lists

An external dynamic list (EDL) is a text file that you host on an external web server. This text file is used so that the firewall can import the following objects:

- IP addresses
- URLs
- Domains

This arrangement allows the firewall to enforce policy based on the entries in the text file list. As you update the list, the firewall dynamically imports the list and enforces policy without the need to make a configuration change or a commit.

The firewall supports the following types of external dynamic lists:

- Predefined IP address
- IP address
- Domain
- URL

You can add a maximum of 30 custom EDLs on your firewall. The EDL list limit is not applicable to Panorama.

Built-In EDLs

An active Threat Prevention license is required to obtain Palo Alto Networks built-in EDLs. These built-in EDLs protect your network against malicious hosts. Built-in EDLs include the following:

- Palo Alto Networks Bulletproof IP Addresses
- Palo Alto Networks High-Risk IP Addresses
- Palo Alto Networks Known Malicious IP Addresses

With the Threat Prevention license, the firewall receives updates for these feeds in content updates. You cannot modify the contents of built-in EDLs.

References

- Formatting Guidelines for an External Dynamic List
  https://docs.paloaltonetworks.com/pan-os/9-1/pan-os-admin/policy/use-an-external-dynamic-list-in-policy/formatting-guidelines-for-an-external-dynamic-list.html
- Built-in External Dynamic Lists
  https://docs.paloaltonetworks.com/pan-os/9-1/pan-os-admin/policy/use-an-external-dynamic-list-in-policy/built-in-edls.html

3.4 Configure application filters and application groups

Application Filters

An administrator can dynamically categorize multiple applications into an application filter based on the specific attributes Category, Subcategory, Tags, Risk, and Characteristic. For example, if you want to allow all audio streaming applications, you could create an application filter that includes the subcategory of audio-streaming, which automatically would add all applications to the filter from the App-ID database that are subcategorized as audio-streaming. The filter then would be added as an application to a Security policy rule. Application filters simplify the process of ensuring that all applications that meet any attribute automatically are added to a Security policy.

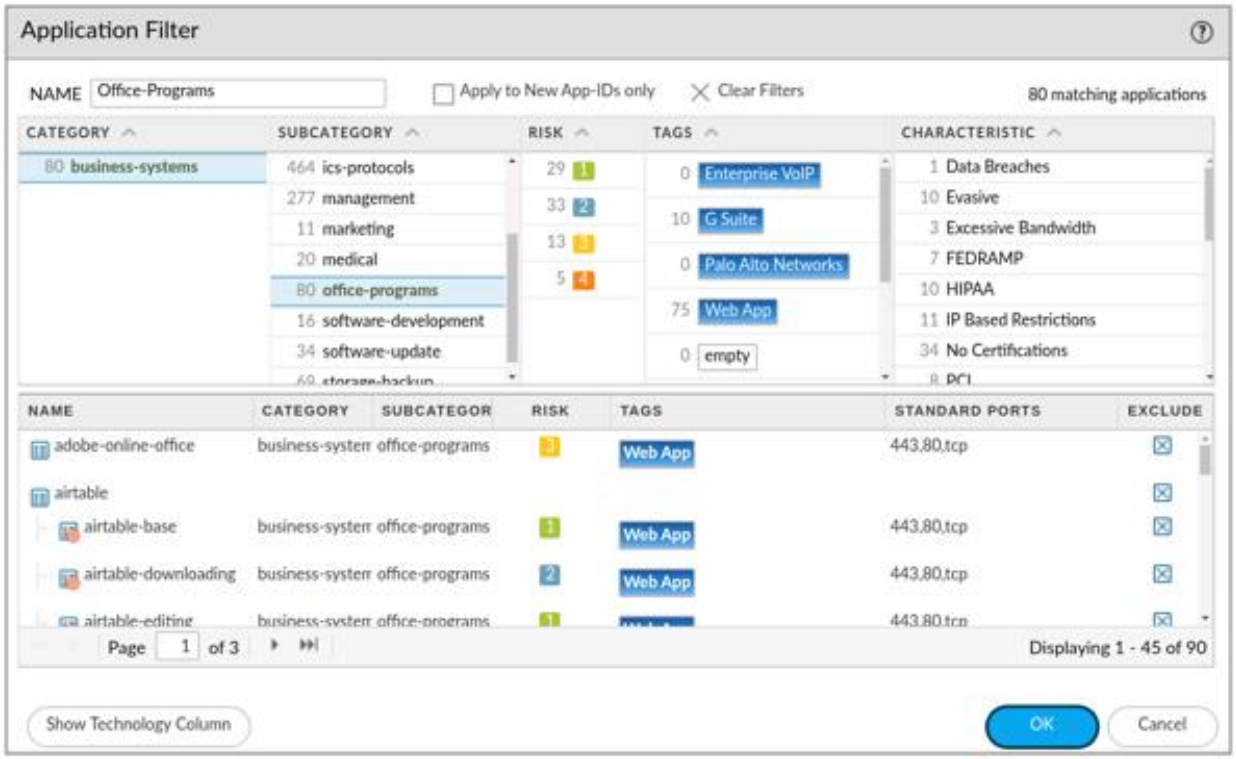

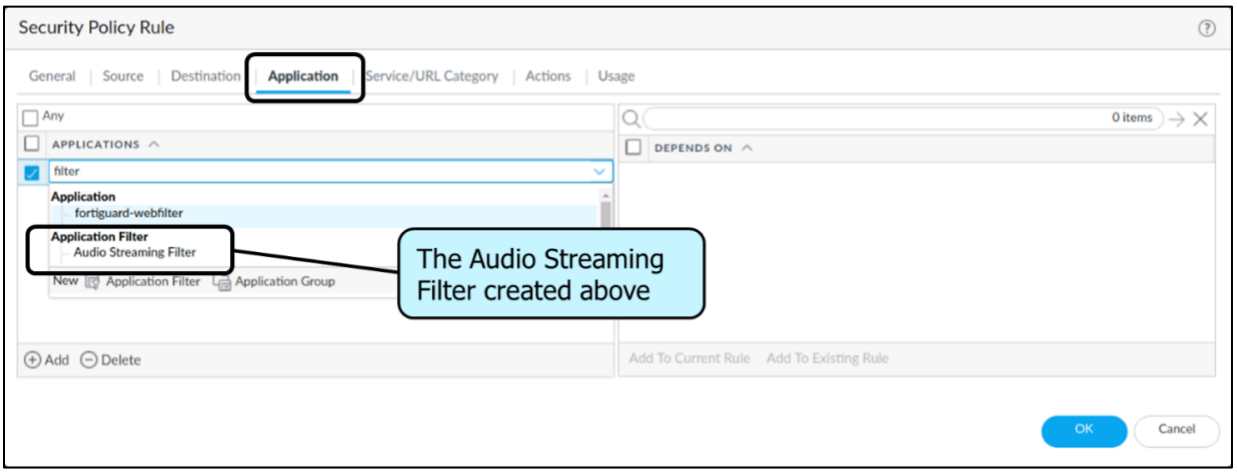

Starting with PAN-OS 9.1, you can configure an application filter to filter for a group of applications based on their assigned application tags. Palo Alto Networks now assigns one or more predefined tags to some applications in the App-ID database. You also can create and assign your own custom tag to an application. You can build an application filter using these tags and then use the application filter in policy rules to control access to the applications. If application tags are updated and they are part of an application filter, then policy could begin to treat such applications differently.

Application Groups

An administrator can manually categorize multiple applications into an application group based on App-ID. This application group then is added to one or more Security policy rules as required, which streamlines firewall administration. Instead of a firewall administrator individually adding different applications into a Security policy, only the application group needs to be added to the policy.

Application groups often are used to simplify security, QoS, and PBF policy rule implementation.

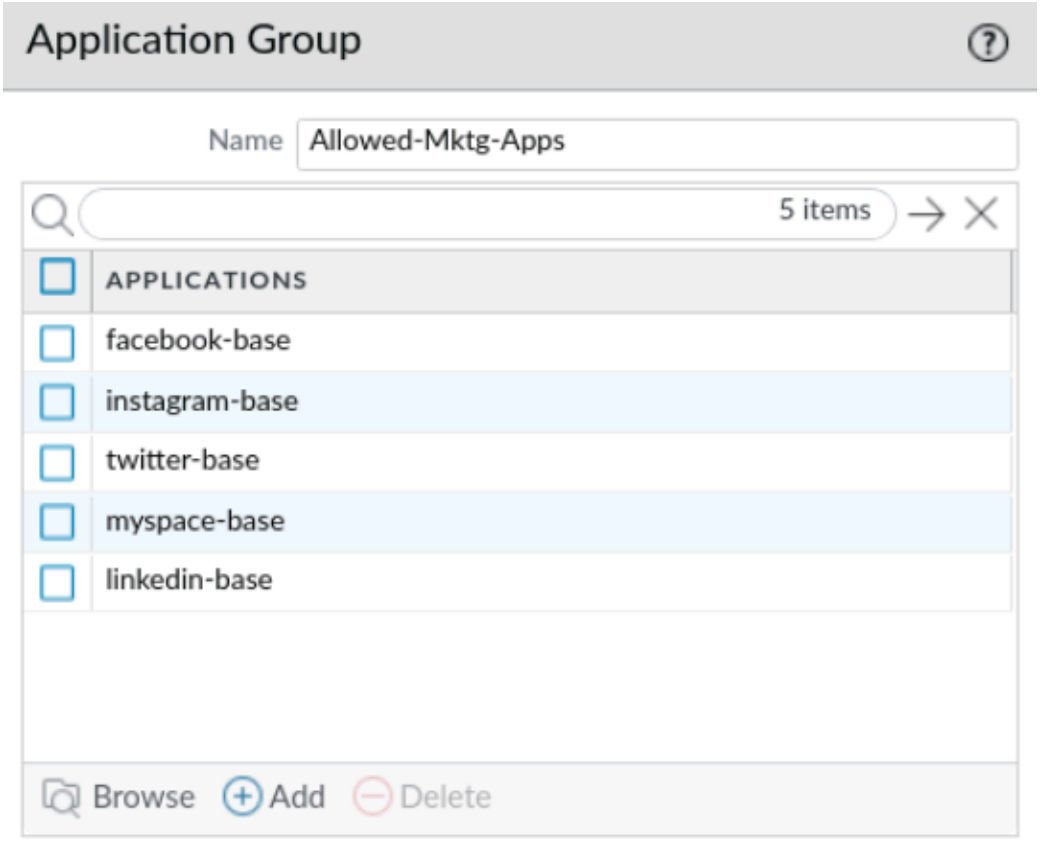

Nesting Application Groups and Filters

An administrator can nest application groups and filters. Multiple applications and multiple application filters can be combined into an application group. One or more application groups then also can be combined into one application group. The final application group then can be added to a Security policy rule.

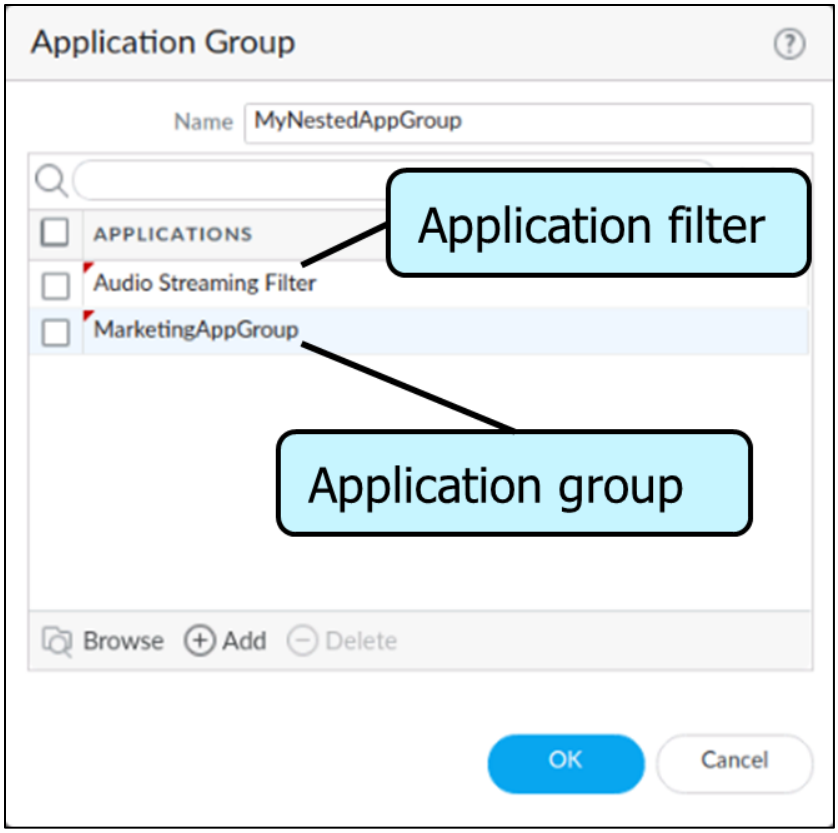

Sample Questions

Q1. What does an application filter enable an administrator to do?

a) manually categorize multiple service filters

b) dynamically categorize multiple service filters

c) dynamically categorize multiple applications

d) manually categorize multiple applications

Q2. Which two items can be added to an application group? (Choose two.)

a) application groups

b) application services

c) application filters

d) application categories

Domain 4 - Policy Evaluation and Management

4.1 Identify the appropriate application-based security policy

Application Shifts

Applications can change during the lifetime of a session. This behavior is called an "application shift." For example, a user types **`www.icloud.com`** into a web browser to access their iCloud email. This initial request goes out as an HTTP request, and the application is recognized as web-browsing. After the HTTP request is completed, the application is changed to icloud-base. After the icloud-base application is processed, the application changes to icloud-mail.

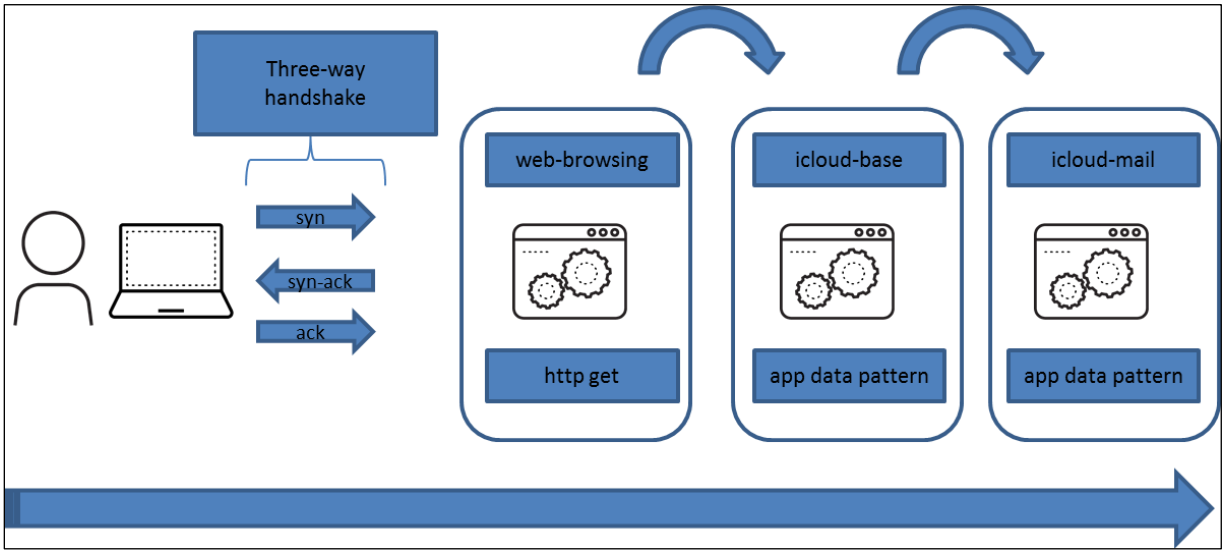

Dependent Applications

Some applications within PAN-OS software are dependent on other applications, which means that if Application#1 is dependent on Application#2, then both Application#1 and Application#2 need to be allowed in the Security policy. For example, icloud-mail is dependent on icloud-base, therefore both applications need to be allowed in the Security policy. Also, icloud-base is dependent on web-browsing, so the web-browsing application also needs to be added to the Security policy. Additional dependent applications are shown in the following figure.

Commit Warnings Due to Missing Dependencies

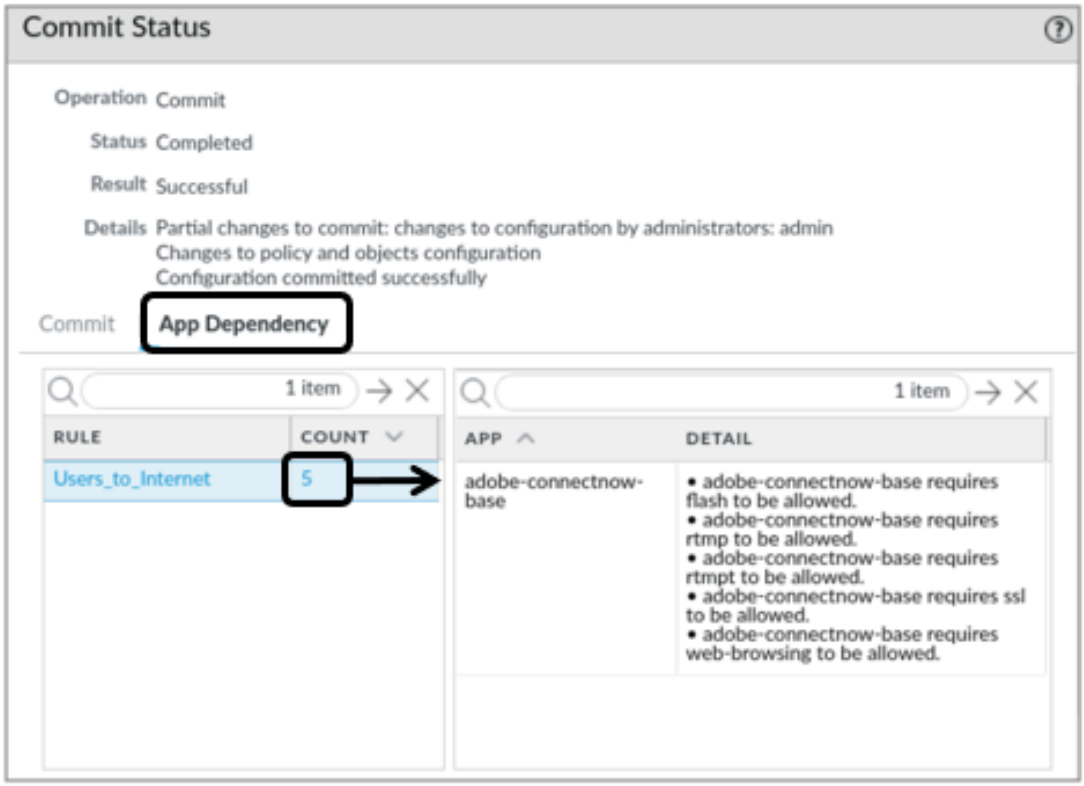

The App Dependency tab does not appear in the Commit Status window if there are no dependency warnings.

Determining Dependent Applications

To determine applications and their dependencies, navigate to **Objects > Applications**:

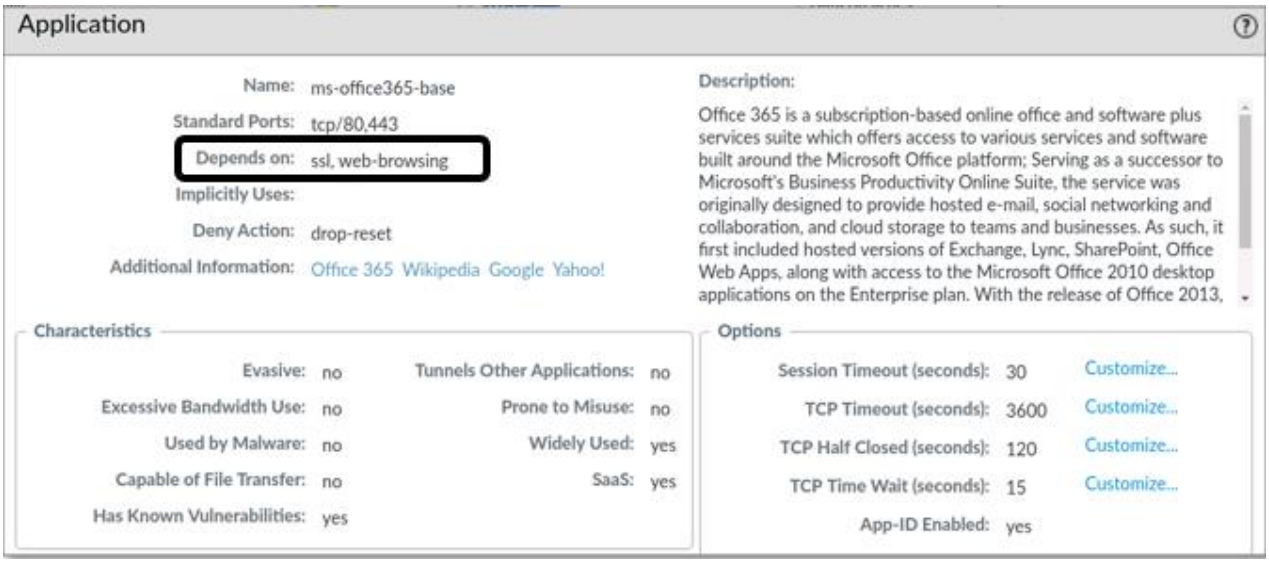

Another method to determine applications and their dependencies is to use the **Depends On** panel on the **Applications** tab while creating or updating a Security policy rule:

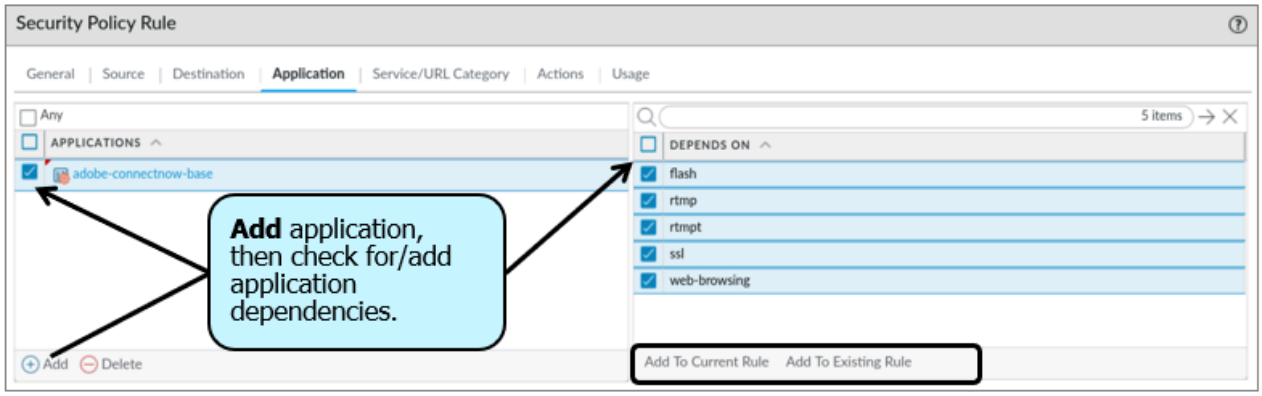

Implicit Applications and Determining Implicit Applications

Some applications such as icloud are dependent on the web-browsing application to be specified in a Security policy. Sometimes you do not have to explicitly allow access to the dependent applications for the traffic to flow because the firewall can determine the dependencies and allow them implicitly. One example is facebook-base. To be able to use facebook-base, you do not have to add ssl or web-browsing to a Security policy.

To determine applications that specifically are used, navigate to **Objects > Applications**:

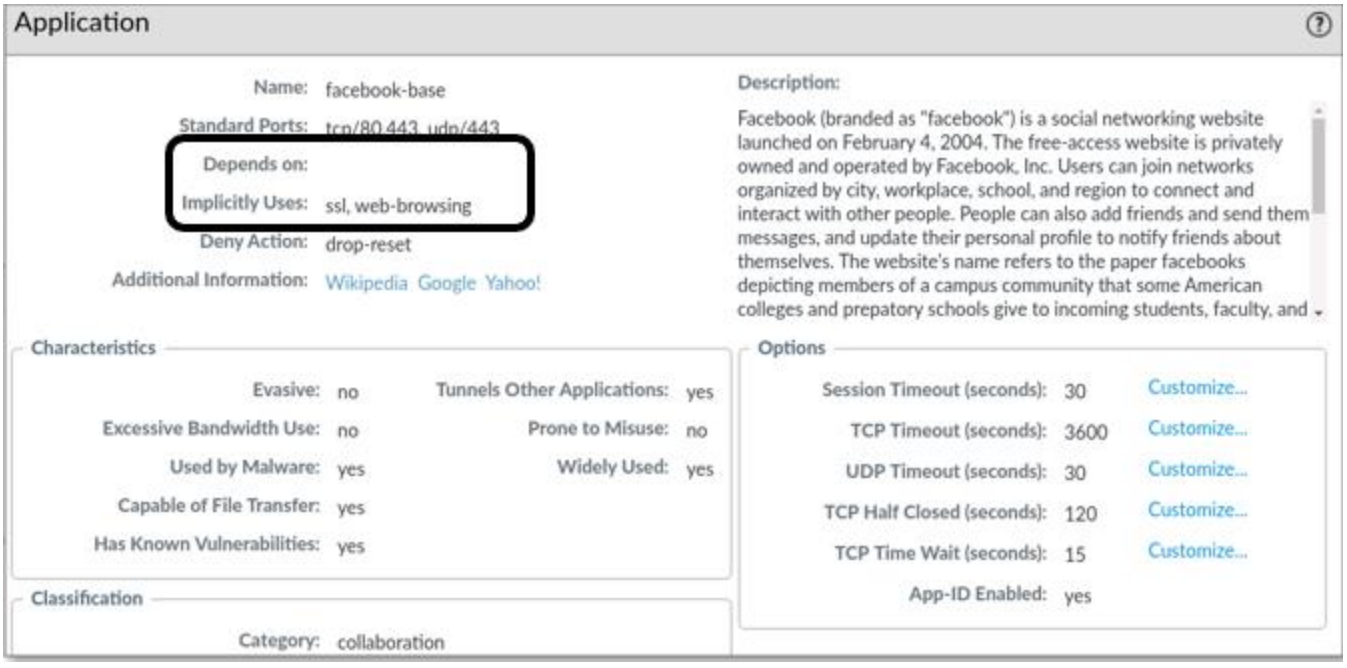

Sample Question

Q1. What are two application dependencies for adobe-connectnow-base? (Choose two.)

a) ssl

b) skype

c) rtmp

d) adobe-base

e) ssh

Dependent Applications

Some applications within PAN-OS software depend on other applications, which means that if Application#1 is dependent on Application#2, then both Application#1 and Application#2 need to be allowed in the Security policy rule for the application to be allowed through the firewall. If applications dependencies are not included in a Security policy rule, then dependency warnings will appear in the **App Dependency** tab of the **Commit Status** window. For example, icloud-mail depends on icloud-base, therefore both applications need to be allowed in the Security policy for icloud-mail to work properly. Also, icloud-base depends on web-browsing, so the web-browsing application also needs to be added to the Security policy. Additional dependent applications are shown in the following figure.

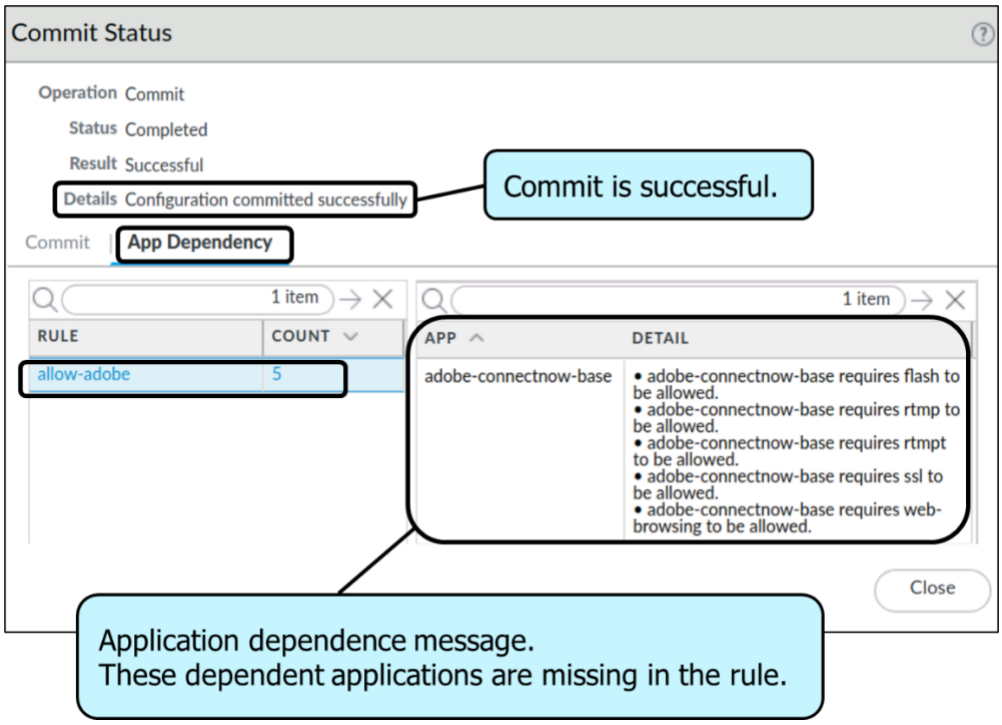

Note: The **App Dependency** tab does not appear in the **Commit Status** window if there are no dependency warnings. Also note that the commit completed successfully even with an application missing a dependent application.

Determining Dependent Applications

You can see application dependencies when you are either creating or updating Security policy rule and when performing commits. When a policy does not include all application dependencies, you can directly access the associated Security policy rule to add the required applications. To determine whether an application has a dependency and what the dependency is before you create a Security policy rule, navigate to **Objects > Applications**. The **Depends on** field will list whether an application has a dependency and which applications need to be included in your Security policy rule:

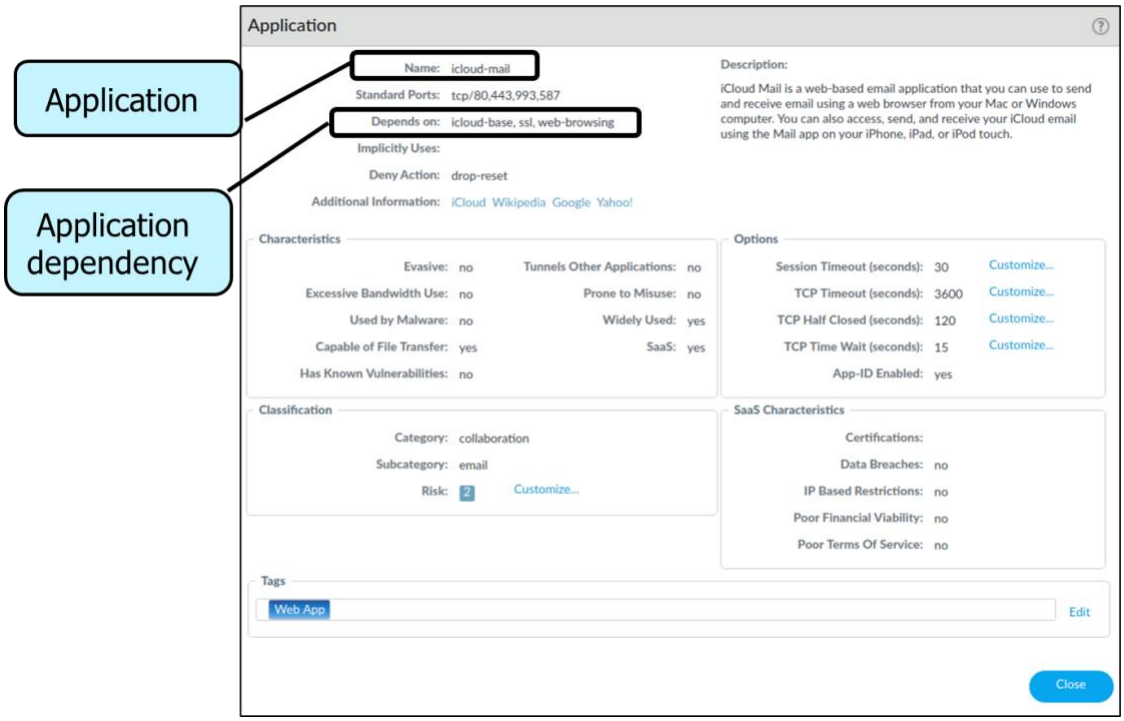

Another method to determine applications and their dependencies while creating or updating a Security policy rule is to use the **Depends On** panel on the **Applications** tab:

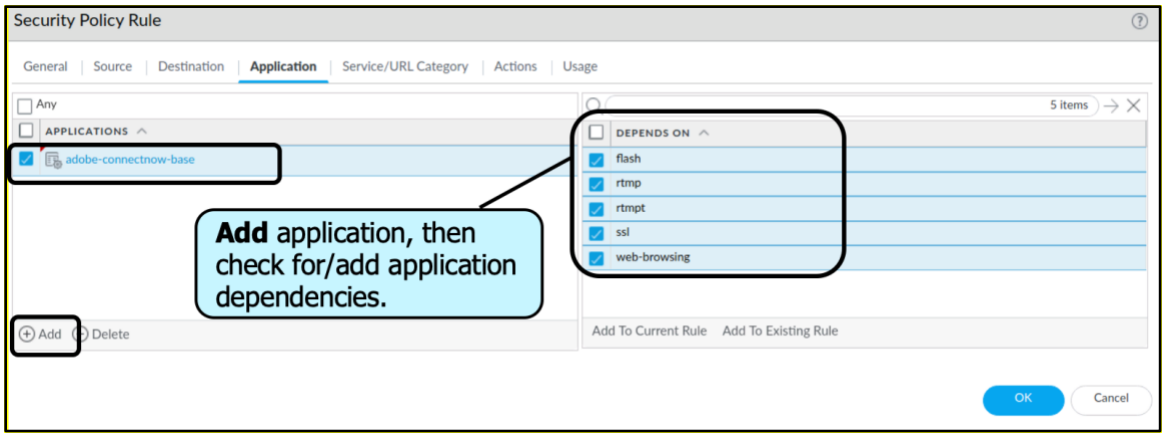

Implicit Applications and Determining Implicit Applications

Some applications such as icloud depend on the web-browsing application to be specified in a Security policy. Sometimes you do not have to explicitly allow access to the dependent applications for the traffic to flow because the firewall can determine the dependencies and allow them implicitly. One example is google-base. To be able to use google-base, you do not have to add ssl or web-browsing to a Security policy.

To determine applications that specifically are used, navigate to **Objects > Applications**. The **Implicitly Uses** field will list the applications that are implicitly allowed as part of the application:

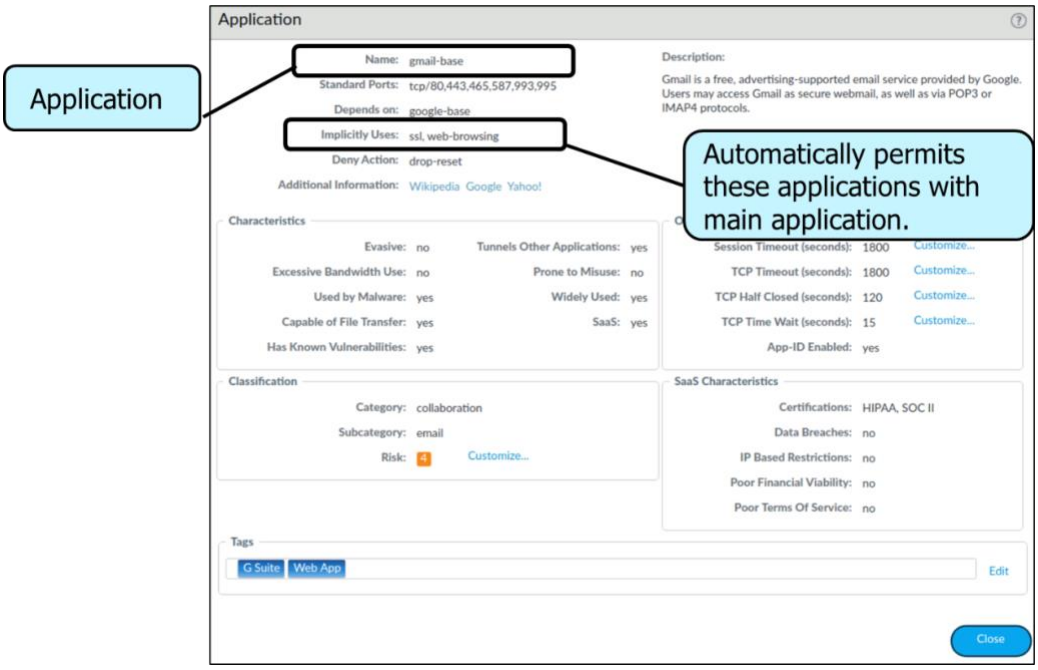

Reference

- Resolve Application Dependencies:

https://docs.paloaltonetworks.com/pan-os/10-0/pan-os-admin/app-id/use-application-objects-in-policy/resolve-application-dependencies.html

4.2 Identify the purpose of specific security rule types

Security Rule Types

Security policies allow you to enforce rules and take action, and they can be as general or as specific as needed. The list of policy rules is compared from the top down against the incoming traffic. The more specific rules must precede the more general ones, because the first rule that matches the traffic is applied.

The default rules apply for traffic that doesn't match any user-defined rules. These default rules are displayed at the bottom of the security rulebase. The default rules are predefined rules that are part of the predefined configuration and are read-only by default; you can override them and change a limited number of settings, including the tags, action (allow or deny), log settings, and security profiles. The names for the two default rules are intrazone-default and interzone-default.

Default Rules: Intrazone Default and Interzone Default

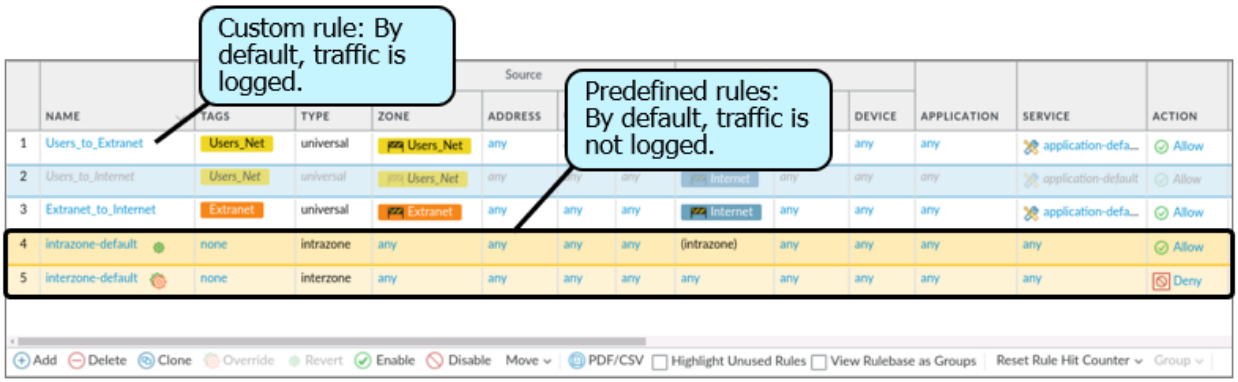

Intrazone Security Policy

Interzone Security Policy

The following table describes the three types of Security policy:

Rule Type	Description
Intrazone • Default rule • Displayed at the bottom of the security rulebase	A Security policy rule allowing traffic within the same zone. Intrazone rule types apply to all matching traffic within the specified source zones (a destination zone cannot be specified for intrazone rules). For example, if the source zone is being set to A and B, the rule would apply to all traffic within zone A and all traffic within zone B, but not to traffic between zones A and B. Traffic logging is not enabled by default. However, best practice is to log the end-of-session traffic.

Rule Type	Description
Interzone • Default rule • Displayed at the bottom of the security rulebase	A security policy rule allowing traffic between two different zones. However, the traffic within the same zone will not be allowed when the policy is created as type Interzone. Interzone rule types apply to all matching traffic between the specified source and destination zones. For example, if the source zone is being set 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 to traffic within zones A, B, or C. Traffic logging is not enabled by default. However, best practice is to log the traffic.
Universal • Exists above the intrazone and interzone security policies	By default, all the traffic destined between two zones, regardless of whether it is from the same zone or different zone. Universal rule types apply to all matching interzone and intrazone traffic in the specified source and destination zones. For example, if a universal rule is being created 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. Traffic logging is enabled by default.

Sample Question

Q1. What are the two default (predefined) Security policy rule types in PAN-OS software? (Choose two.)

a) Universal

b) Interzone

c) Intrazone

d) Extrazone

Q2. True or false. Because the first rule that matches the traffic is applied, the more specific rules must follow the more general ones.

a) true

b) false

Q3. Which statement is true?

a) For Intrazone traffic, traffic logging is enabled by default.

b) For Interzone traffic, traffic logging is enabled by default.

c) For Universal traffic, traffic logging is enabled by default.

d) For any rule type, traffic logging is enabled by default.

4.3 Identify and configure Security policy match conditions, actions, and logging options

Implicit and Explicit Rules

Two implicit (predefined) Security policy rules come with the PAN-OS software: intrazone and interzone. The intrazone Security policy rule allows traffic within a zone by default. The interzone Security policy does not allow traffic between zones by default. These two predefined Security policy rule types reside at the bottom of the security rulebase set and are processed after all other preceding Security policy rules are processed. All preceding security rules must be explicitly defined by an administrator. Note that traffic is not logged by default for the predefined rules and that traffic is logged by default for explicitly defined rules. Best practice is to log for all Security policy rules, whether implicit or explicit.

A shadow rule warning indicates that a broader rule matching the criteria is configured above a more specific rule.

The following screenshot shows that no traffic ever will match the second rule, which specifically allows skype and dropbox, because all applications already have been allowed by the first rule. Rule 2's "skype" shadows rule 3's "skype."

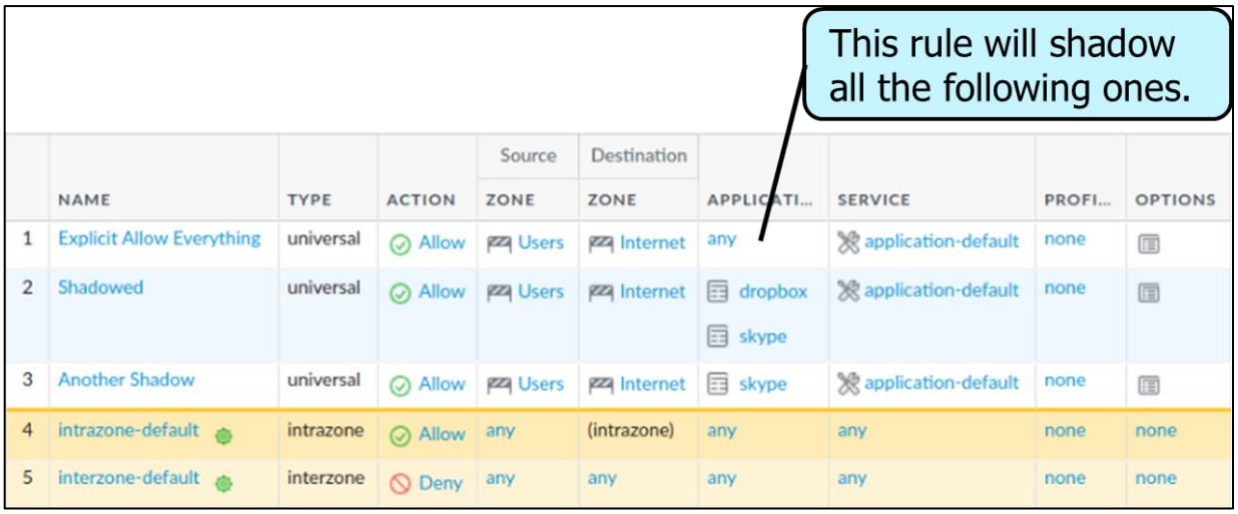

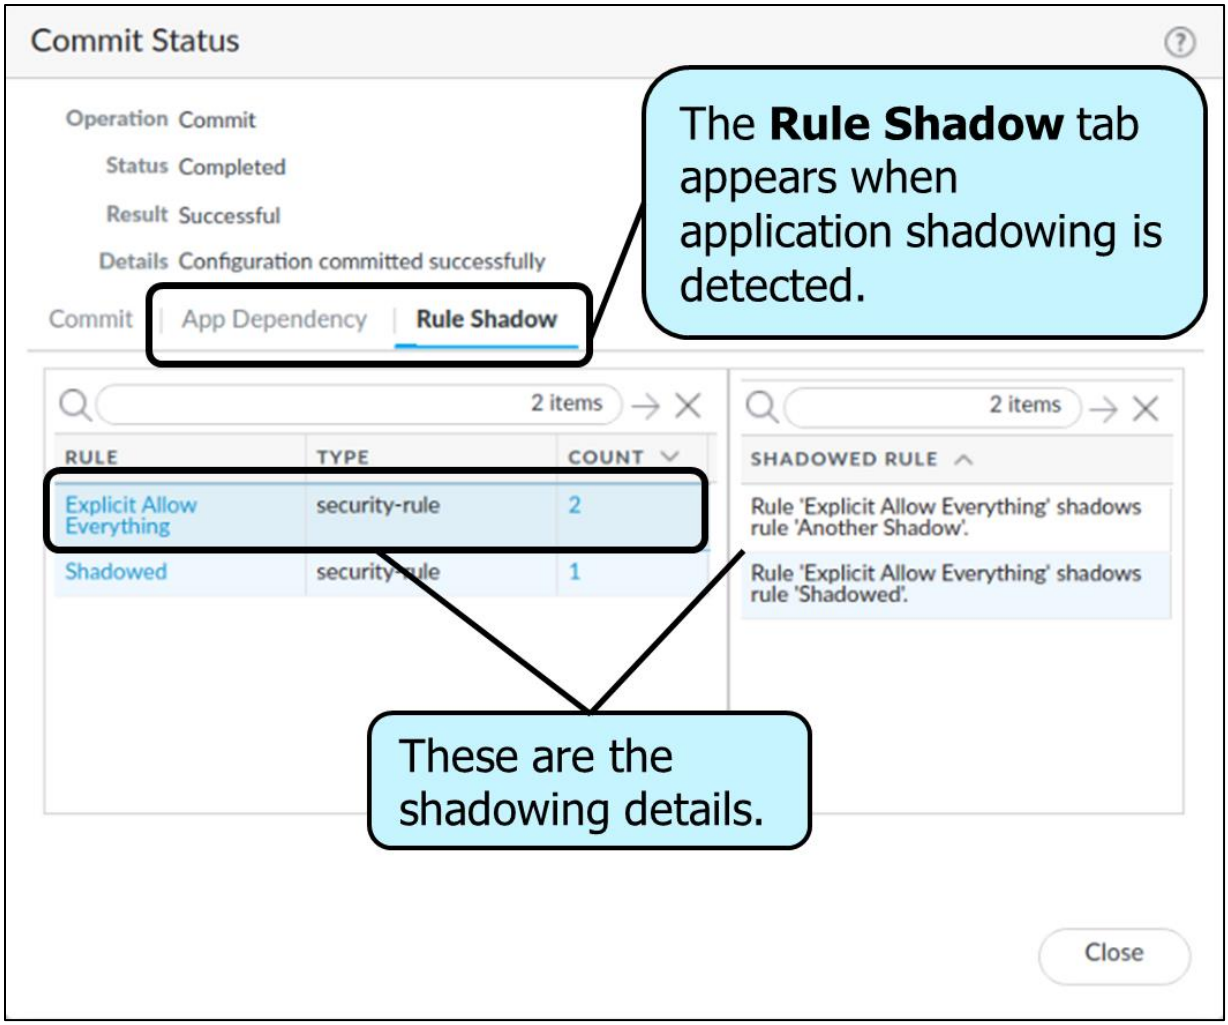

Security Rule Hit Count

The PAN-OS software enables you to monitor hit count. The three components of **Rule Usage** are as follows:

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

In the following screenshot, note that the hit counts have not incremented because this example has no live traffic:

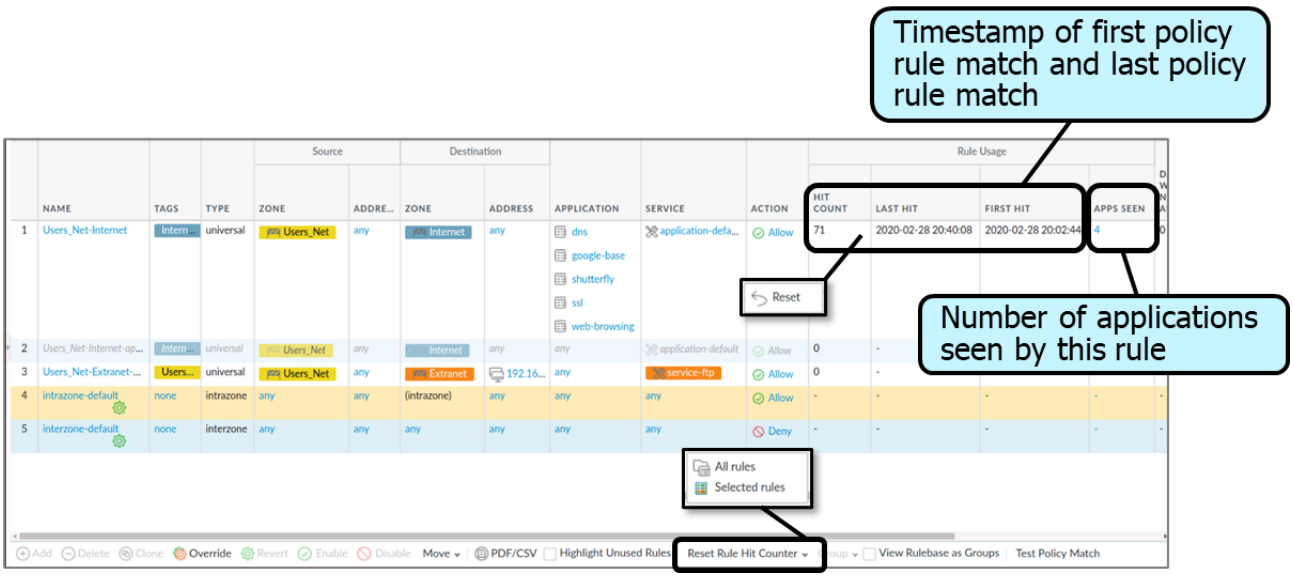

Sample Question

Q1. What are the two default (predefined) Security policy rule types in PAN-OS software? (Choose two.)

a) Universal

b) Interzone

c) Intrazone

d) Extrazone

Q2. What will be the result of one or more occurrences of shadowing?

a) a failed commit

b) an invalid configuration

c) a warning

d) an alarm window

Q3. Which type of Security policy rules most often exist above the two predefined security policies?

a) Intrazone

b) Interzone

c) Universal

d) Global

Application Properties

All applications in the App-ID database are defined by six properties:

Property	Definition
Category	Used to generate the Top Ten Application Categories chart within the ACC and is available for filtering.
Subcategory	Also used to generate the Top Ten Application Categories chart within the ACC and is available for filtering.
Technology	Technology most closely associated with the application.
Parent App	Specify a parent application for this application. This setting applies when a session matches both the parent and the custom applications; however, the custom application is reported because it is more specific.
Risk	A relative risk rating from 1 to 5, with 5 being the most risky.
Characteristics	Identifies some application property or behavior, like certified for FEDRAMP, or can be used for evasion, or can use excessive bandwidth, and so on.

Application Characteristics

All applications in the App-ID database are defined by six characteristics. The names of the characteristics are self-explanatory.

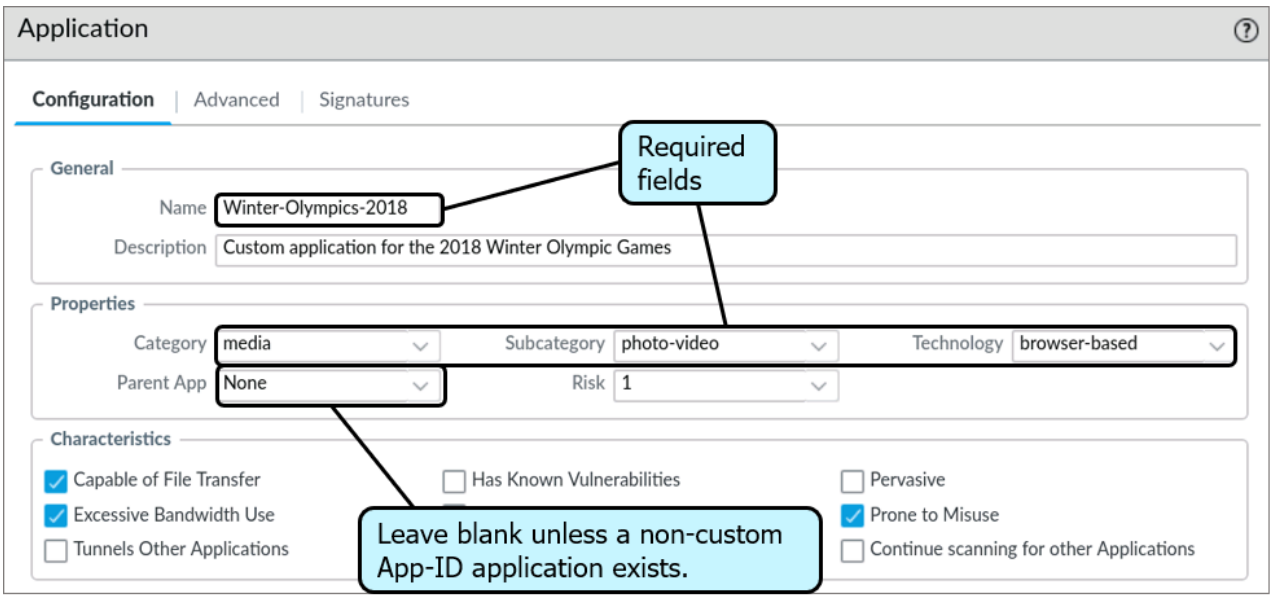

Application Timeouts

Item	Definition
Timeout	Number of seconds before an idle application flow is terminated. A zero indicates that the default timeout of the application will be used. This value is used for protocols other than TCP and UDP in all cases, and for TCP and UDP timeouts when the TCP timeout and UDP timeout are not specified.
TCP Timeout	Number of seconds before an idle TCP application flow is terminated. A zero indicates that the default timeout of the application is used.
UDP Timeout	Number of seconds before an idle UDP application flow is terminated. A zero indicates that the default timeout of the application is used.
TCP Half Closed	Maximum length of time that a session remains in the session table between reception of the first FIN and reception of the second FIN or RST. If the timer expires, the session is closed.
TCP Time Wait	Maximum length of time that a session remains in the session table after the second FIN or RST is received. If the timer expires, the session is closed. If this time is not configured at the application level, the global setting is used (range is 1 to 600 seconds). If this value is configured at the application level, it overrides the global TCP Time Wait setting.

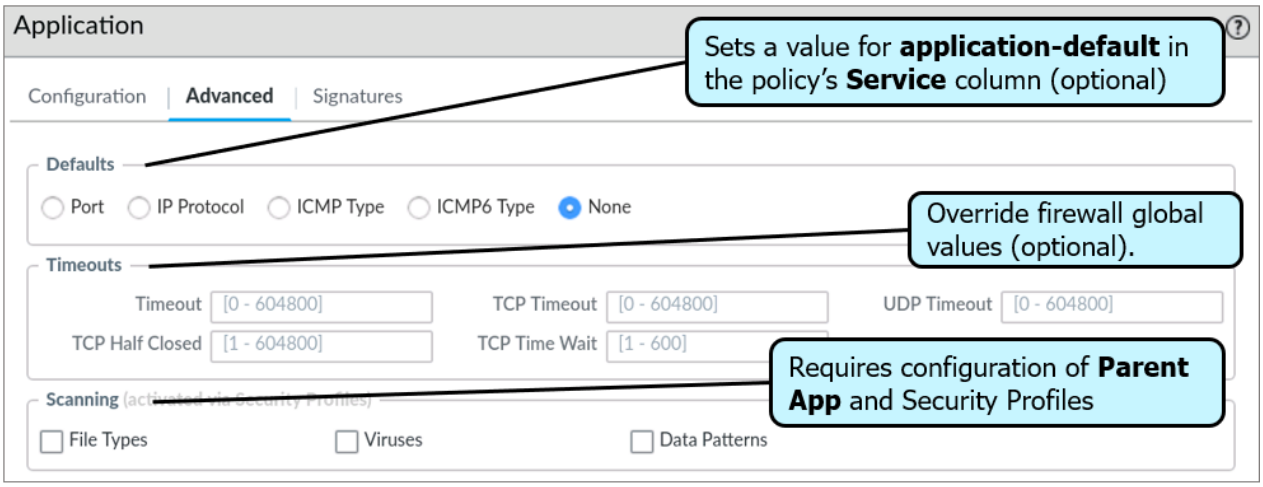

Sample Questions

Q4. What does the TCP Half Closed setting mean?

a) maximum length of time that a session remains in the session table between reception of the first FIN and reception of the third FIN or RST

b) minimum length of time that a session remains in the session table between reception of the first FIN and reception of the second FIN or RST

c) maximum length of time that a session remains in the session table between reception of the first FIN and reception of the second FIN or RST

d) minimum length of time that a session remains in the session table between reception of the first FIN and reception of the third FIN or RST.

Q5. What are two application characteristics? (Choose two.)

a) stateful

b) excessive bandwidth use

c) intensive

d) evasive

4.4 Identify and implement the proper NAT policy

NAT Types

The two basic types of NAT are source NAT (SNAT) and destination NAT (DNAT).

SNAT is used to replace the original source IP address in a packet. A typical scenario for SNAT is when a packet is originated from within a company's network and then is forwarded out to the internet. The original source IP address usually is an RFC 1918 IP address that is not routable within the internet.

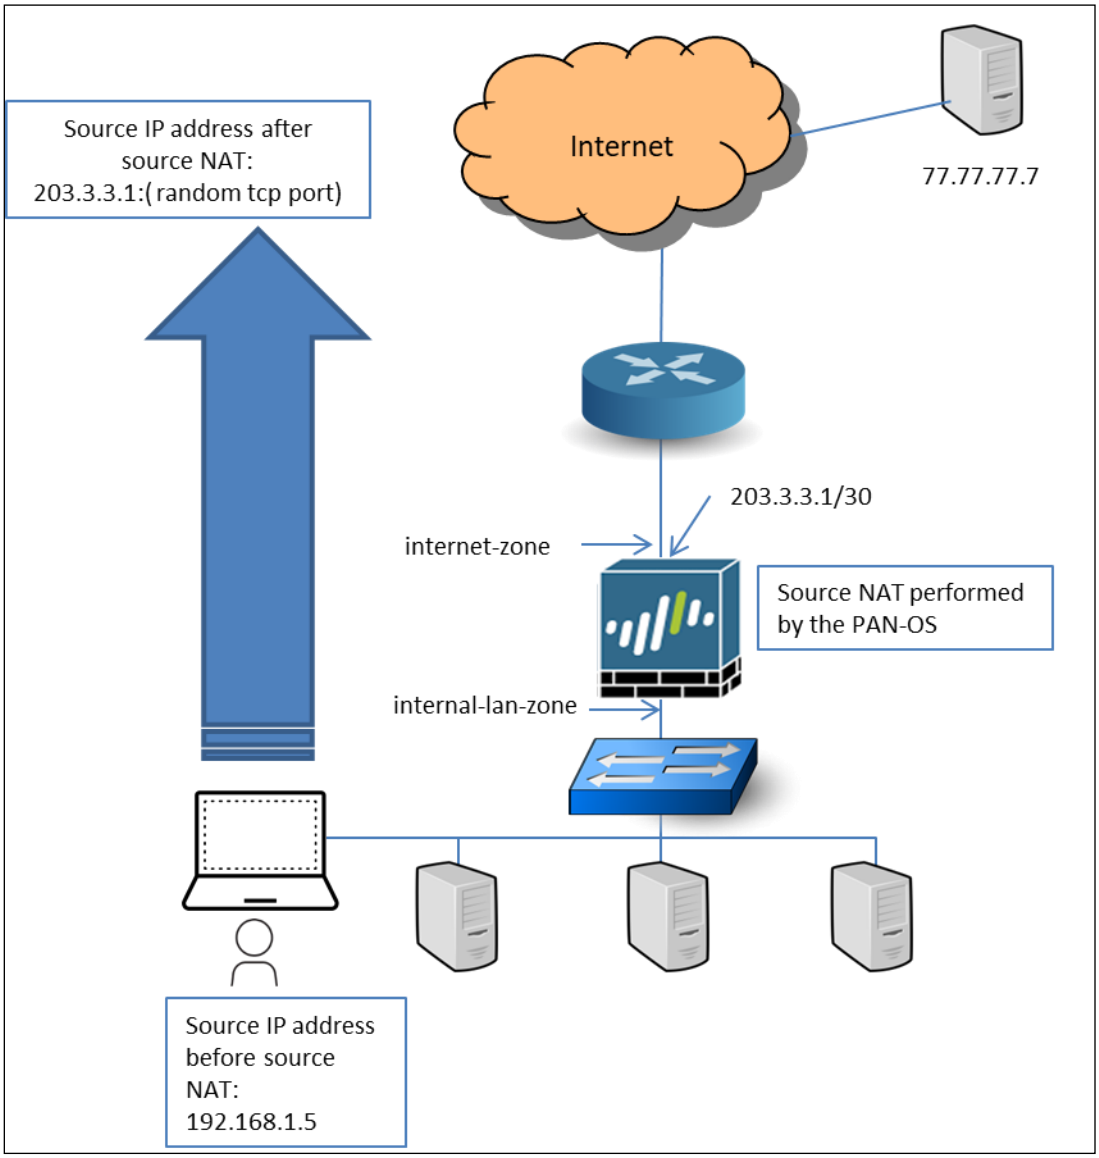

DNAT is used to replace the original destination IP address in a packet. A typical scenario for DNAT is when a packet is originated from the internet and then is forwarded to a company's internal network. The original destination IP is routable within the internet. When the packet arrives at the firewall, the routable IP address is replaced with the real IP address of the destination device (usually an RFC 1918 IP address) and then is forwarded to the destination device.

DNAT can be used in other scenarios such as when subnets overlap.

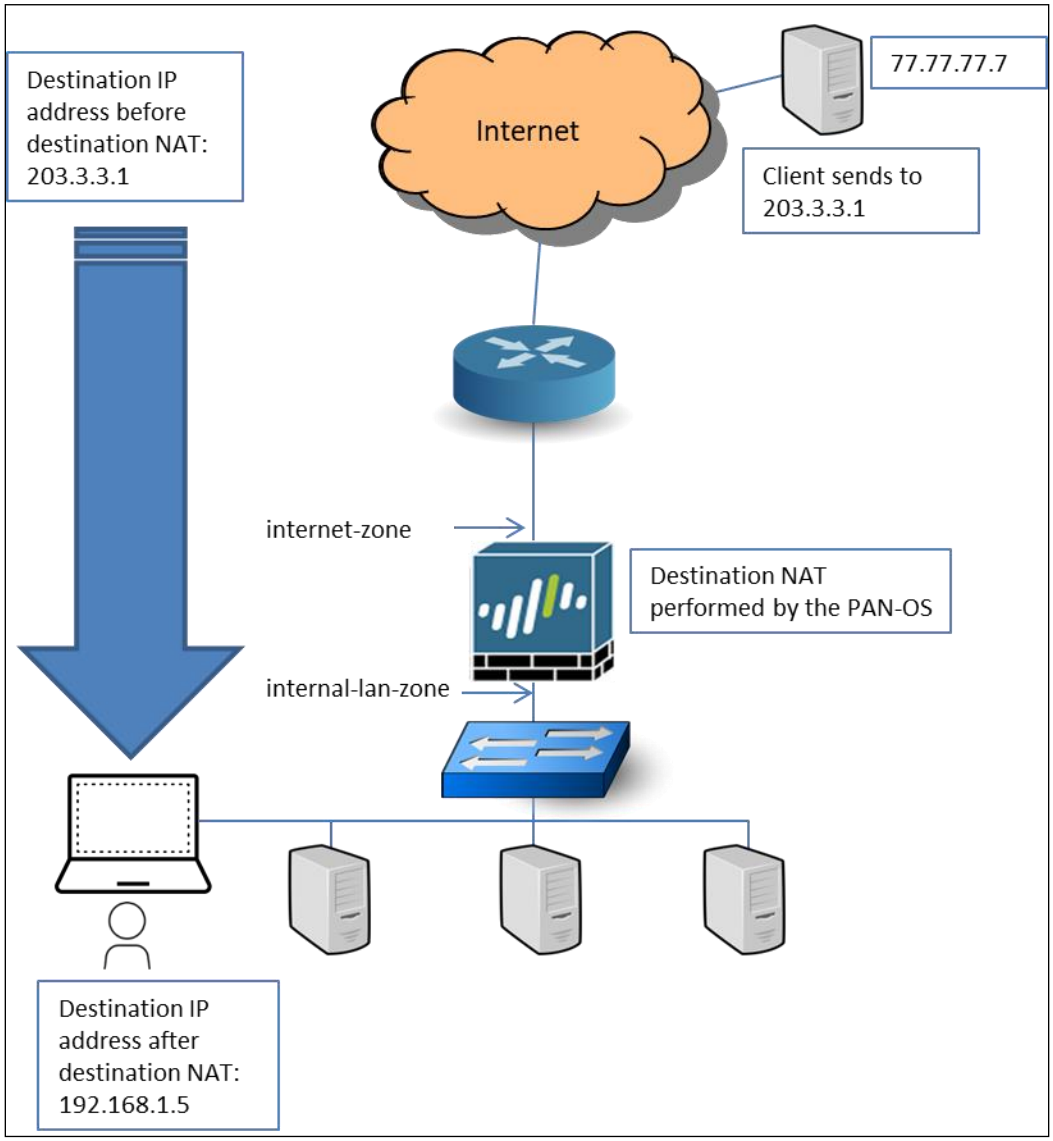

Source NAT Types

The following table describes the three source NAT types: static IP, dynamic IP, and dynamic IP and port:

Source NAT Type	Description
Static IP	The same address always is used for the translation and the port is unchanged. For example, if the source range is 192.168.0.1 – 192.168.0.10, and the translation range is 10.0.0.1 – 10.0.0.10, address 192.168.0.2 always is translated to 10.0.0.2. The address range usually is limited. This concept applies if only a host /32 IP address is used.
Dynamic IP	The original source IP address translates to the next available address in the specified range but the port number remains unchanged. Up to 32,000 consecutive IP address are supported. A dynamic IP pool can contain multiple subnets, so you can translate your internal network addresses to two or more separate public subnets.
Dynamic IP and port	This is the most commonly used source NAT type. Address selection is based on a hash of the source IP address. For a given source IP address, the firewall uses the same translated source address for all sessions.

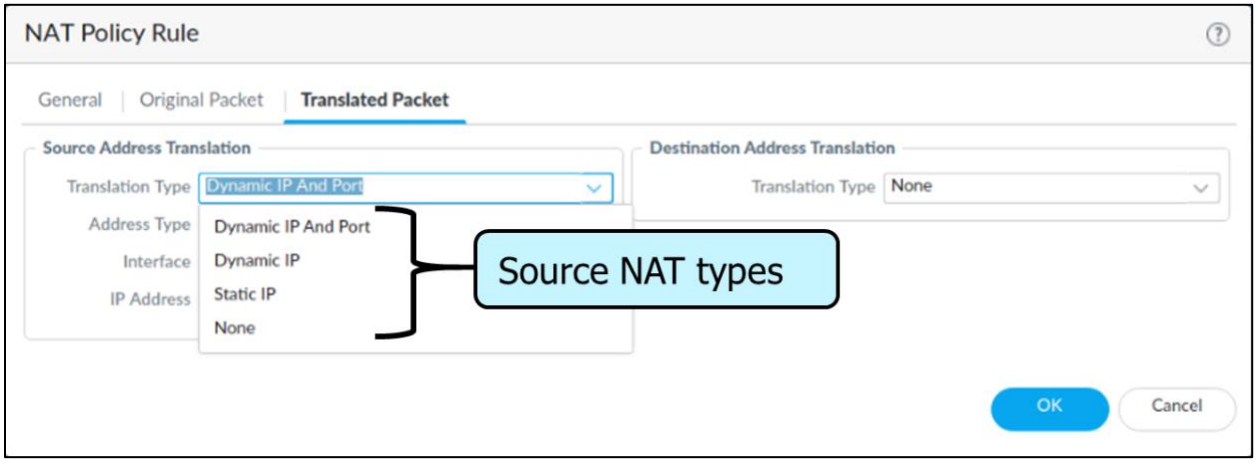

Source NAT and Security Policies

A Security policy rule requires a source IP, destination IP, source zone, and destination zone. If you use an IP address in a Security policy rule, you must add the IP address value that existed before NAT was implemented, which is called the pre-NAT IP. After the IP address is translated (post-NAT IP), determine the zone where the post-NAT IP address would exist. This post-NAT zone is used in the Security policy rule.

A simple way to remember how to configure Security policy rules where NAT was implemented is to memorize the following: “pre-NAT IP; post-NAT zone.”

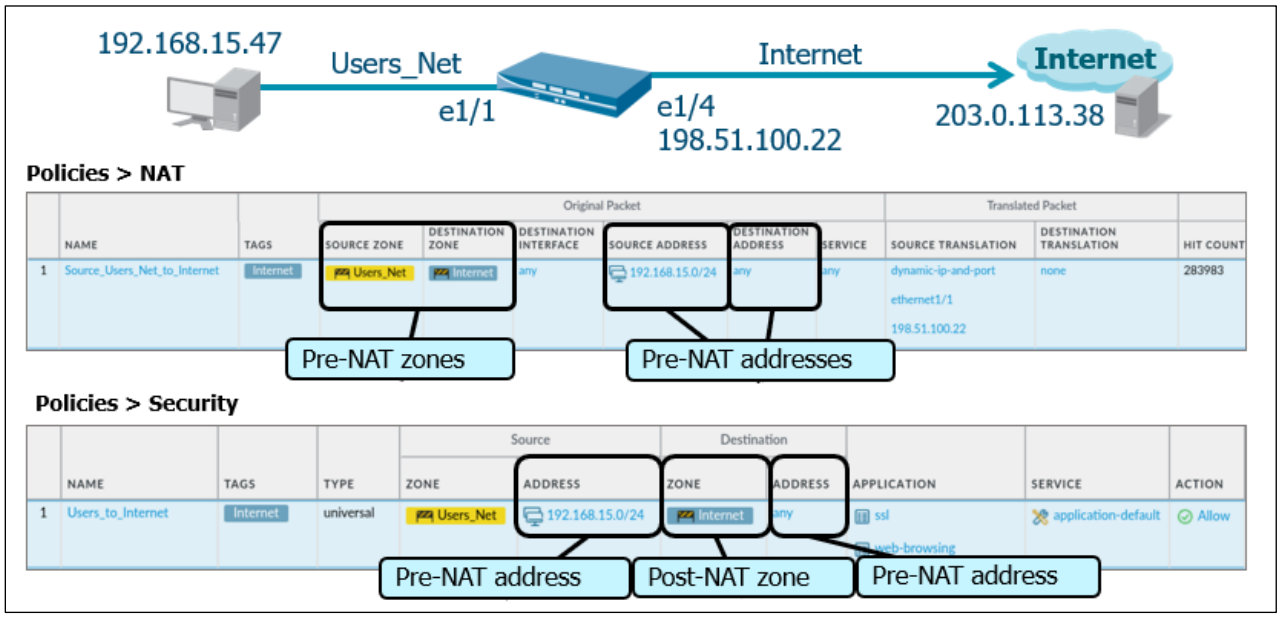

Configuring Source NAT

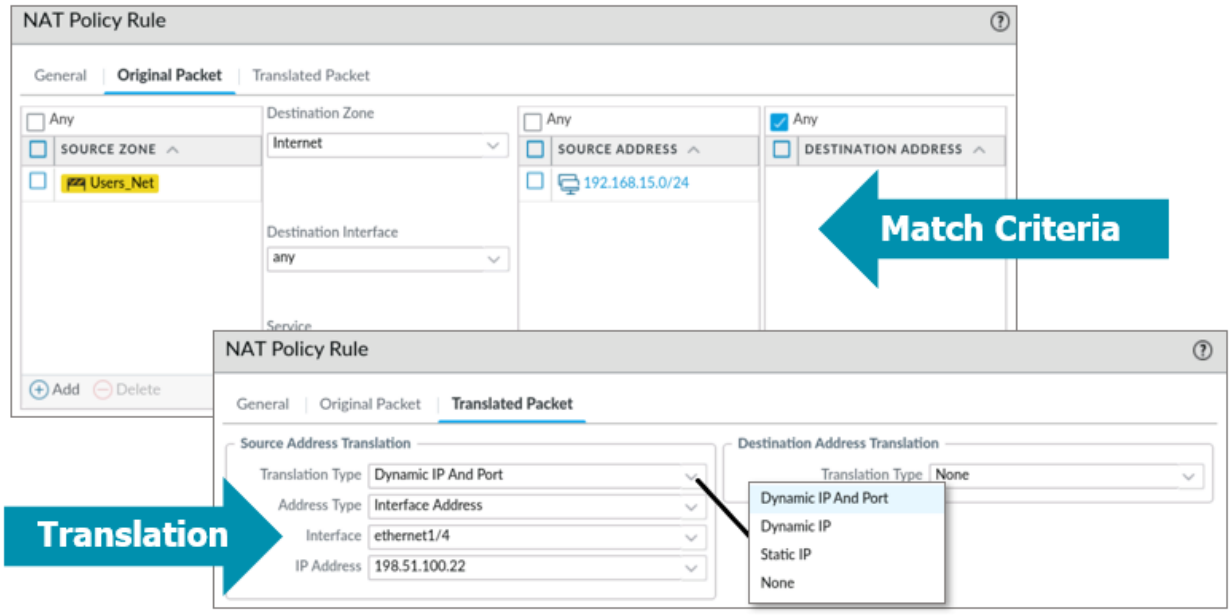

Configuring Bidirectional Source NAT

For static translations, bidirectional NAT enables the firewall to create a corresponding translation in the opposite direction of the translation you configure. If you are configuring static source NAT, bidirectional NAT enables you to eliminate the need to create an additional NAT policy rule for the incoming traffic.

If you enable bidirectional translation, you must ensure that you have Security policy rules in place to control the traffic in both directions. If there are no such rules, the bidirectional feature allows packets to be translated automatically in both directions.

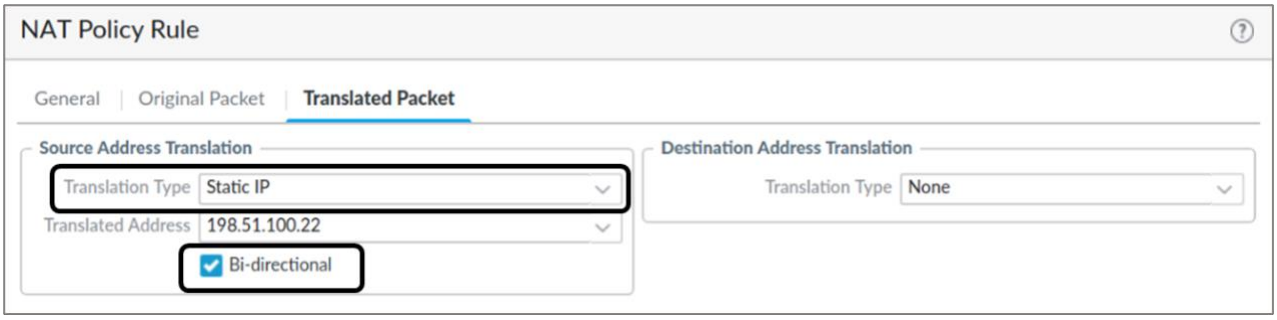

DIPP NAT Oversubscription

The DIPP **NAT Oversubscription Rate** is the number of times that the same translated IP address and port pair can be used concurrently. Reduction of the oversubscription rate will decrease the number of source device translations but will provide higher NAT rule capacities. Oversubscription assumes that the destination is different in each translation.

Platform Default turns off oversubscription, whereby the default rate of the firewall model applies:

- 1x: means no oversubscription, where each IP address and port pair can be used only one time
- 2x: oversubscribed two times
- 4x: oversubscribed three times
- 8x: oversubscribed eight times

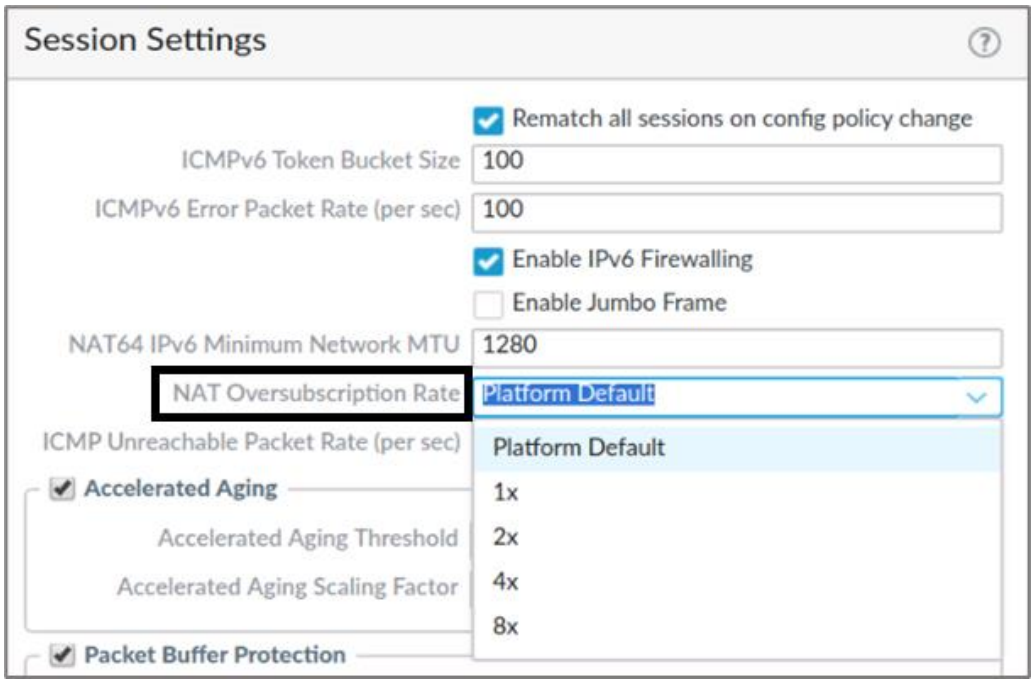

Destination NAT Types

Destination NAT (DNAT) typically is used to allow an external client to initiate access to an internal host such as a web server. The two types of destination NAT are as follows:

Destination NAT Type	Description
Static	You can set the translated address as an IP address or range of IP addresses and a translated port number (1 – 65,535), to which the original destination address and port number are translated. If the **Translated Port** field is blank, the destination port is not changed.
Dynamic IP (with session distribution)	You can enter a translated address that is an FQDN, an address object, or an address group from which the firewall selects the translated address. If the DNS server returns more than one address for an FQDN or if the address object or address group translates into more than one IP address, the firewall distributes sessions among those addresses using the specified session distribution method.

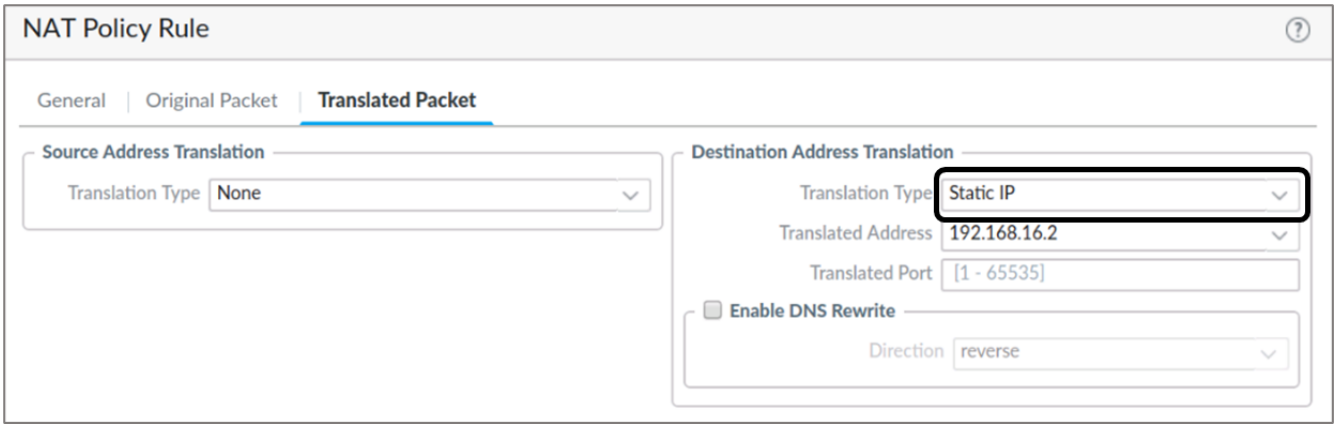

Destination NAT and Security Policies

A Security policy rule requires a source IP, destination IP, source zone, and destination zone. If you use an IP address in a Security policy rule, you must add the IP address value that existed before NAT was implemented, which is called the pre-NAT IP. After the destination IP address is translated (post-NAT IP), determine the zone where the post-NAT IP address would exist. This post-NAT zone is used in the Security policy rule.

A simple way to remember how to configure security policy rules where NAT was implemented is to memorize the following: “pre-NAT IP; post-NAT zone.”

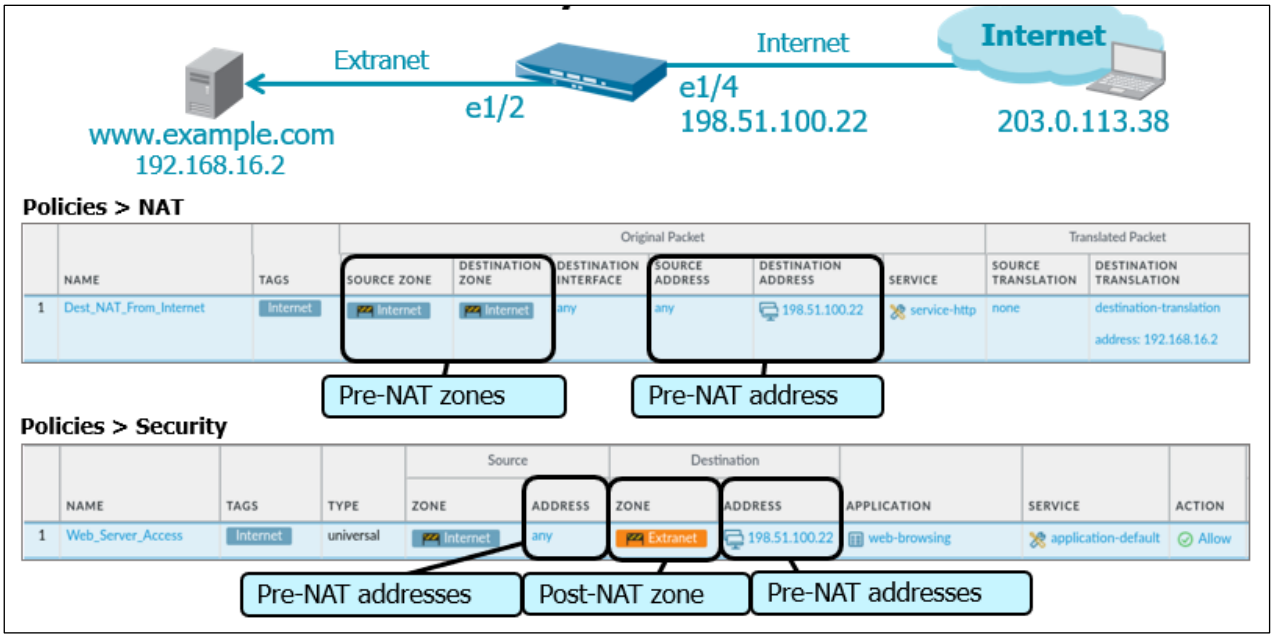

Configuring Destination NAT

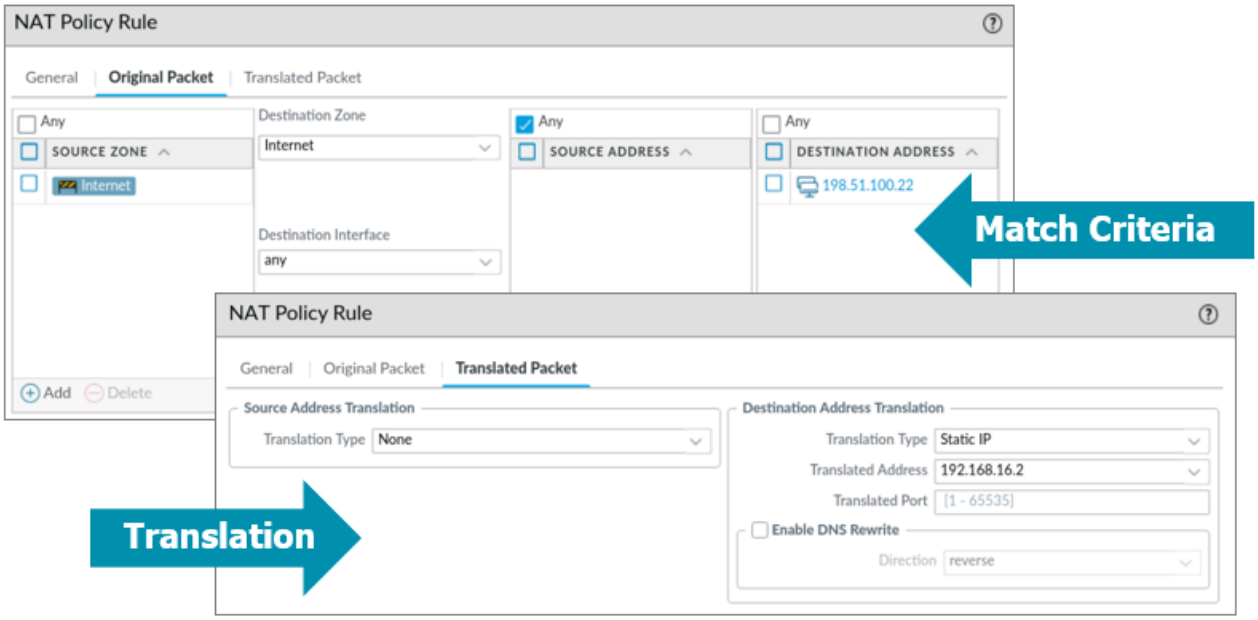

Configuring Dynamic IP Address Support for DNAT

You can enter a translated address that is an FQDN, an address object, or an address group from which the firewall selects the translated address. If the DNS server returns more than one address for an FQDN or if the address object or address group translates into more than one IP address, the firewall distributes sessions among those addresses using the specified session distribution method.

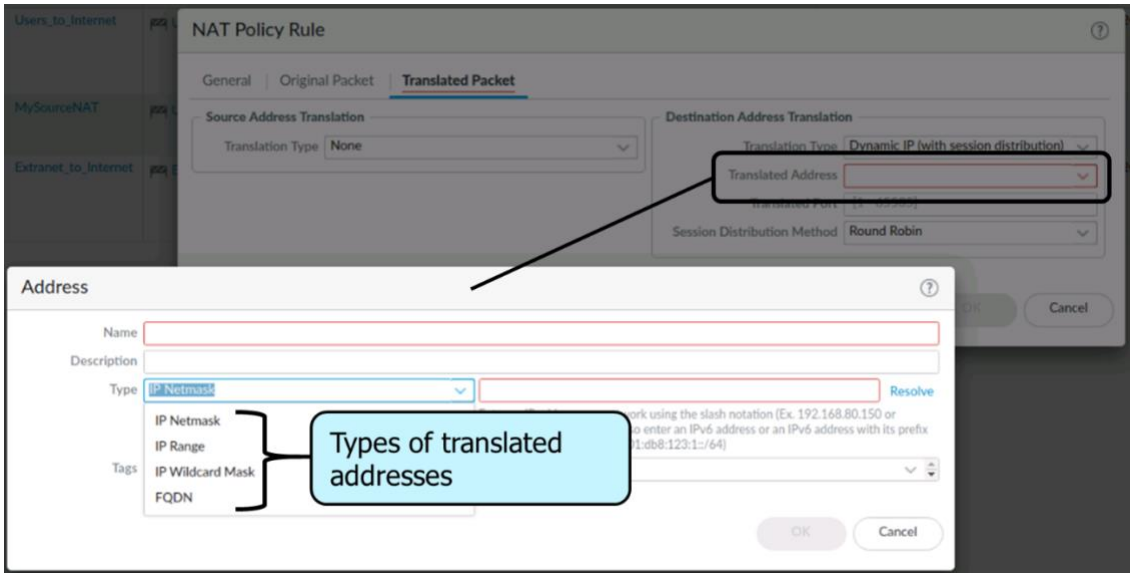

Configuring Destination NAT Port Forwarding

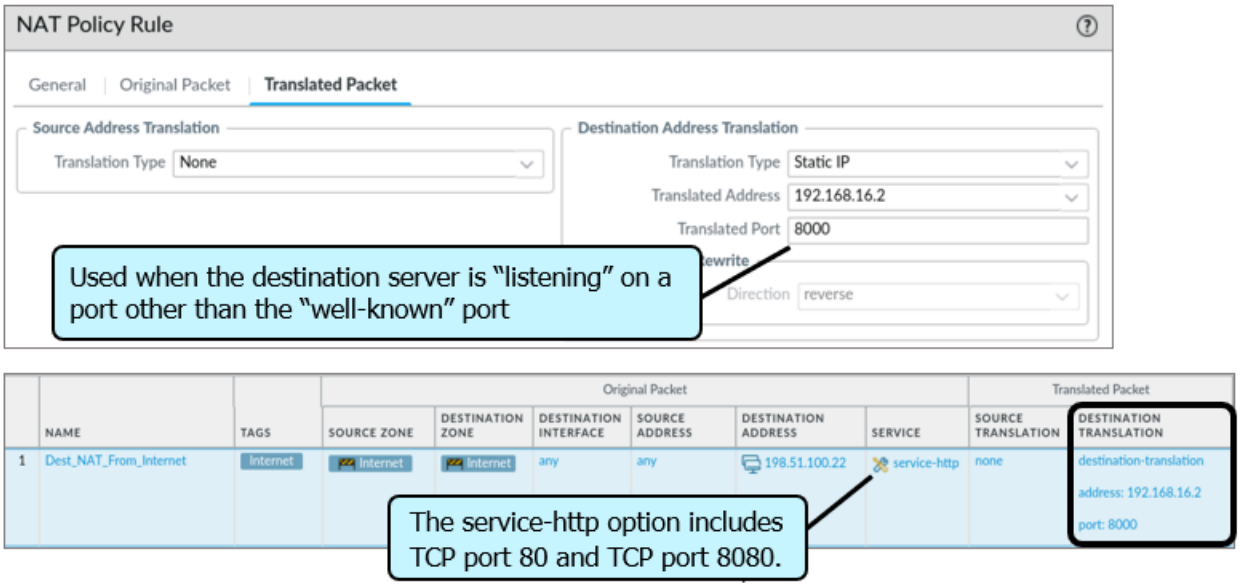

Sample Question

Q1. What are two source NAT types? (Choose two.)

a) universal

b) static

c) dynamic

d) extrazone

Q2. Which phrase is a simple way to remember how to configure Security policy rules where NAT was implemented?

a) post-NAT IP, pre-NAT zone

b) post-NAT IP, post-NAT zone

c) pre-NAT IP, post-NAT zone

d) pre-NAT IP, pre-NAT zone

Q3. What are two types of destination NAT? (Choose two.)

a) dynamic IP (with session distribution)

b) DIPP

c) global

d) static

Q4. What are two possible values for DIPP NAT oversubscription? (Choose two.)

a) 1x

b) 4x

c) 16x

d) 32x

Q4. Which statement is true regarding bidirectional NAT?

a) For static translations, bidirectional NAT enables the firewall to create a corresponding translation in the opposite direction of the translation you configure.

b) For static translations, bidirectional NAT enables the firewall to create a corresponding translation in the same direction of the translation you configure.

c) For dynamic translations, bidirectional NAT enables the firewall to create a corresponding translation in the opposite direction of the translation you configure.

d) For dynamic translations, bidirectional NAT enables the firewall to create a corresponding translation in the same direction of the translation you configure.

4.5 Identify the tools available to optimize Security policies

Policy Optimizer provides a simple workflow to migrate your legacy Security policy rulebase to an App-ID-based rulebase, which improves your security by reducing the attack surface and offering visibility into applications so you can safely enable them. Policy Optimizer identifies port-based rules so you can convert them to application-based whitelist rules or add applications from a port-based rule to an existing application-based rule without compromising application availability. It also identifies over-provisioned App-ID-based rules (App-ID rules configured with unused applications). Policy Optimizer helps you prioritize which port-based rules to migrate first, identify application-based rules that allow applications you do not use, and analyze rule usage characteristics such as hit count.

Conversion of port-based rules to application-based rules improves your security posture because you select the applications you want to whitelist and deny all other applications. You therefore eliminate unwanted and potentially malicious traffic from your network. The combination of restricting application traffic to its default ports (set the **Service** to **application-default**) and conversion to application-based rules also prevents evasive applications from running on non-standard ports.

Using the Policy Optimizer

Begin by identifying existing port-based rules.

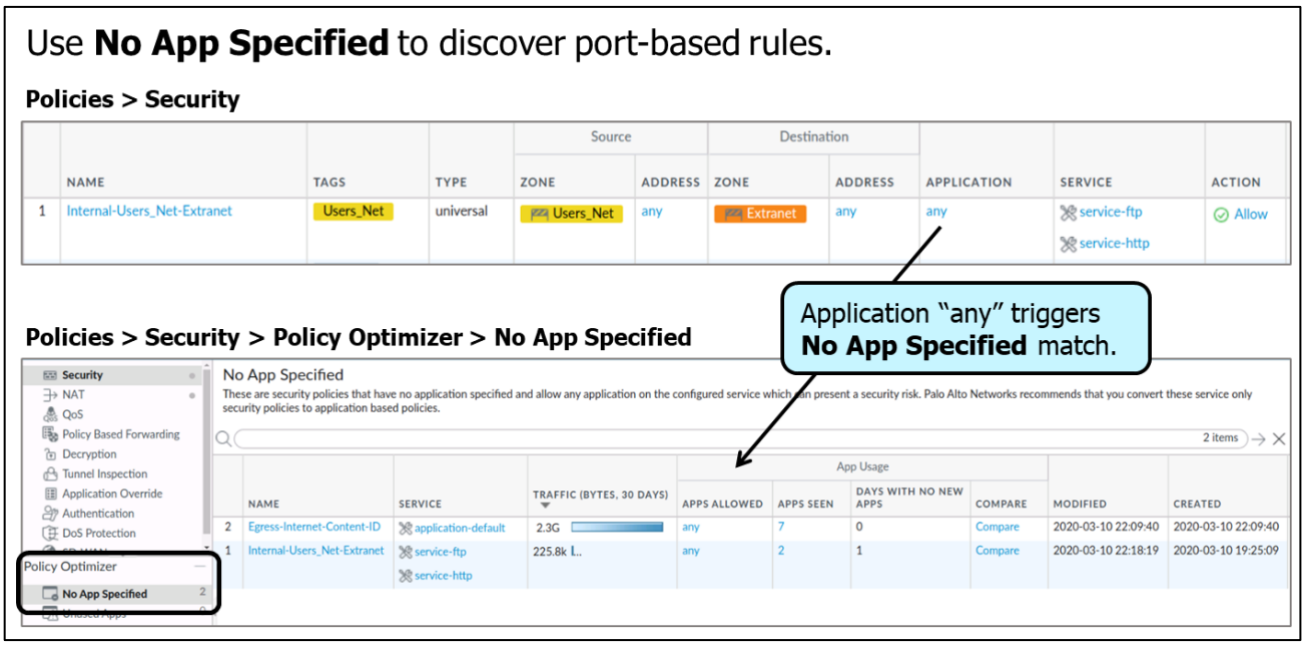

Use the displayed information to determine which applications were seen in a specific rule. Click **Compare** to display the list and application usage information:

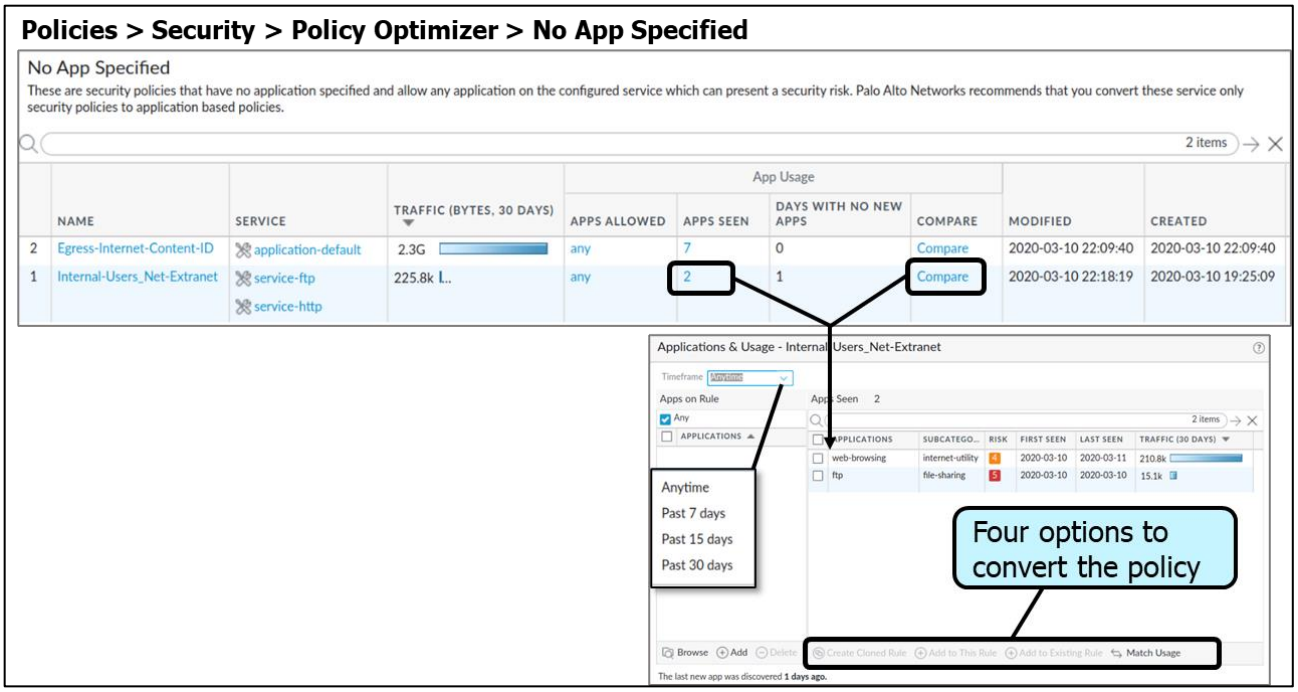

You can add the discovered App-IDs to your Security policy rules using one of the four selections at the bottom of the **Application & Usage** display. Select a listed **Application** by checking the box to make the options available.

- Create Cloned Rule: Creates a duplicate of the Security policy rule being examined, adding the selected App-ID
- Add to This Rule: Adds the selected **Application**(s) to the existing Security policy rule being evaluated
- Add to Existing Rule: Adds the selected **Application**(s)to an existing Security policy rule of your choice
- Match Usage: Replaces the selected Security policy rule with a new, App-ID-based version in the same Security policy rule list position

The Policy Optimizer also can display existing App-IDs in Security policy rules that have not been used, and identify Security policy rules that have not matched traffic at 30 day, 60 day, and never intervals.

Sample Question

Q1. The Policy Optimizer does not analyze which statistics?

a) applications allowed through port-based Security policy rules

b) the usage of existing App-IDs in Security policy rules

c) which users matched security policies

d) existing Security policy rule App-IDs that have not matched processed traffic

e) days since the latest new application discovery in a port-based Security policy rule

App-ID Updates and Impact

Firewall administrators must be careful before they install any App-ID updates because some applications might have changed since the last App-ID update (content update). For example, an application that previously was categorized under web-browsing now might be categorized under its own unique App-ID. Categorization of applications into more specific applications enables more granularity and control of applications within Security policy rules. Because the new App-ID no longer will be categorized as web-browsing, no Security policy rule now will contain this new App-ID. Consequently, the new App-ID will be blocked.

You can minimize this risk by using the **Disable new apps in content** update feature. New updates will be downloaded and installed according to the schedule, but they will be disabled until they are manually enabled. Be aware that this action will force you to track disabled applications, which increases your administrative burden. You may want to examine the effect of any new applications on your Security policy and make any required policy updates without disabling new application signatures.

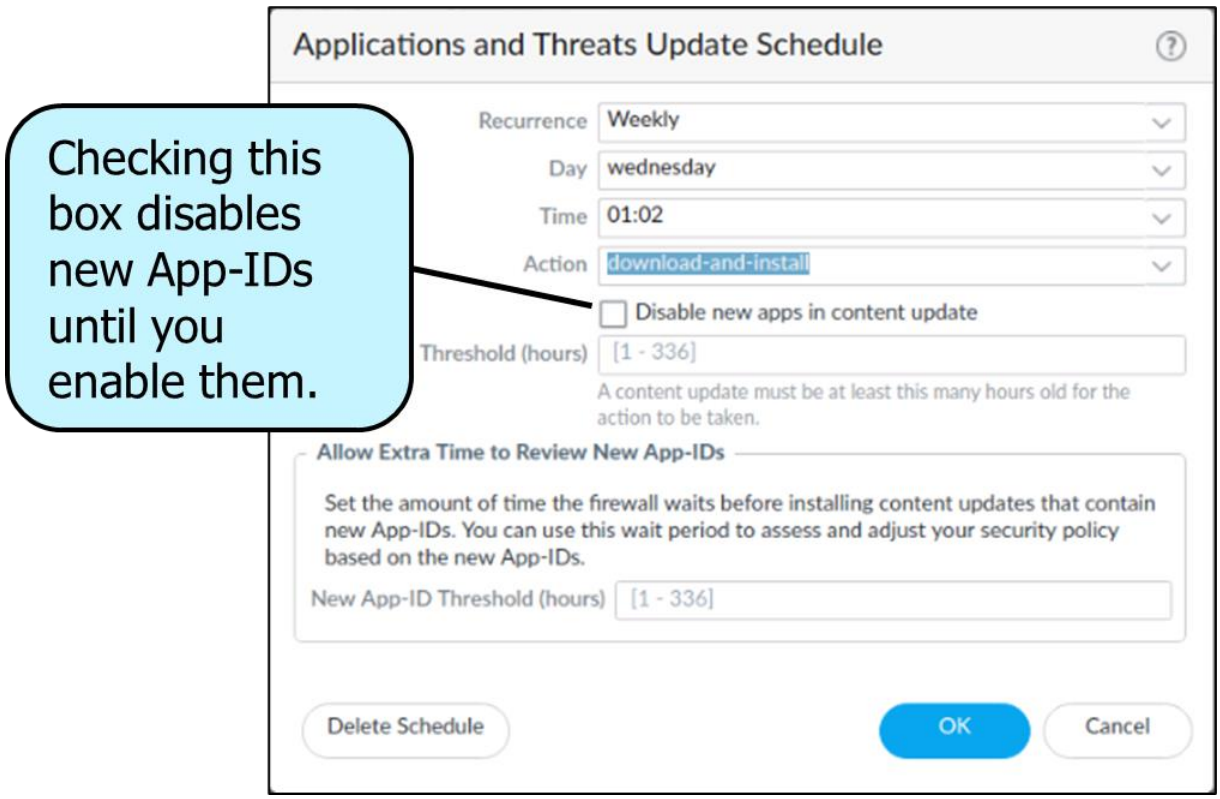

Content Update Absorption

To see the applications that have been modified since the last content release, select **Review Apps** in the **Action** column. The screen will display details about the modified application.

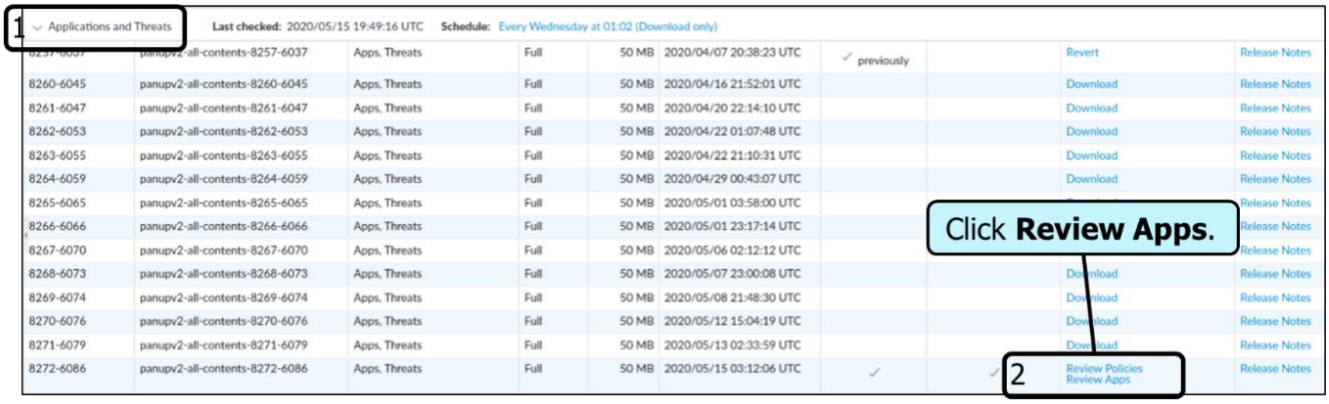

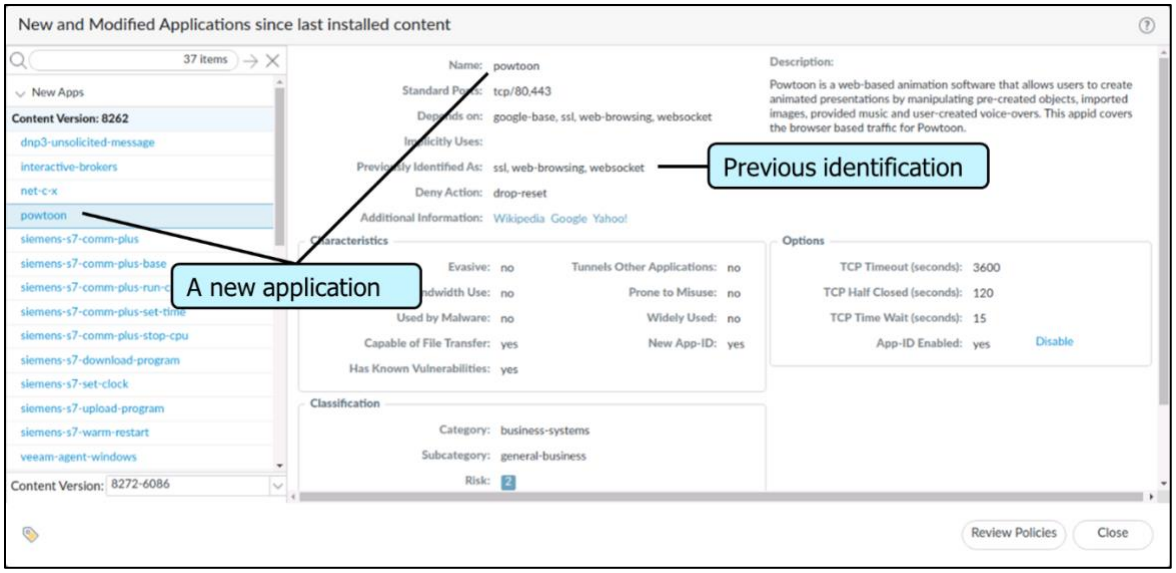

Select **Review Policies** to display the Security policy rules that might enforce traffic differently after the application is modified:

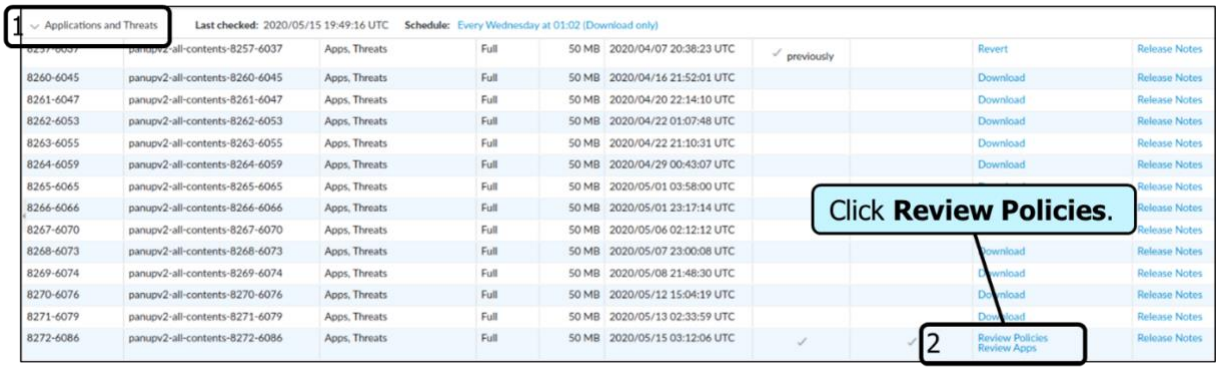

Always review content release notes for the list of newly identified and modified applications and threat signatures that the content release introduces. Content release notes also describe how the update might impact existing Security policy enforcement and provide recommendations about how you can modify your Security policy to best leverage what's new. Installation of new App-IDs included in a content release version sometimes can cause a change in policy enforcement for the application that now is uniquely identified.

Sample Questions

Q2. Which column in the **Applications and Threats** screen includes the options **Review Apps** and **Review Policies**?

a) Features

b) Type

c) Version

d) Action

Q3. Which link can you select in the web interface to minimize the risk using of installing new App-ID updates?

a) Enable new apps in content

b) Disable new apps in App-ID database

c) Disable new apps in content update

d) Enable new apps in App-ID database

Application Tags

Starting with the release of PAN-OS 9.1, Palo Alto Networks adds predefined tags to many applications listed in the App-ID database. The predefined tags are assigned to applications based on the application's characteristics. For example, web-based applications are assigned the tag *Web App* and VoIP applications are assigned the tag *Enterprise VoIP*. Predefined tags are updated and maintained by the *Applications and Threats* dynamic updates.

You can view application tags in the web interface by browsing to **Objects > Applications** and then opening an application's details window (see the following image). You can create custom tags using **Objects > Tags** and then assign your custom tag to an application. To assign a custom tag to an application, use the **Edit** link in the application's details window, as shown in the following screenshot:

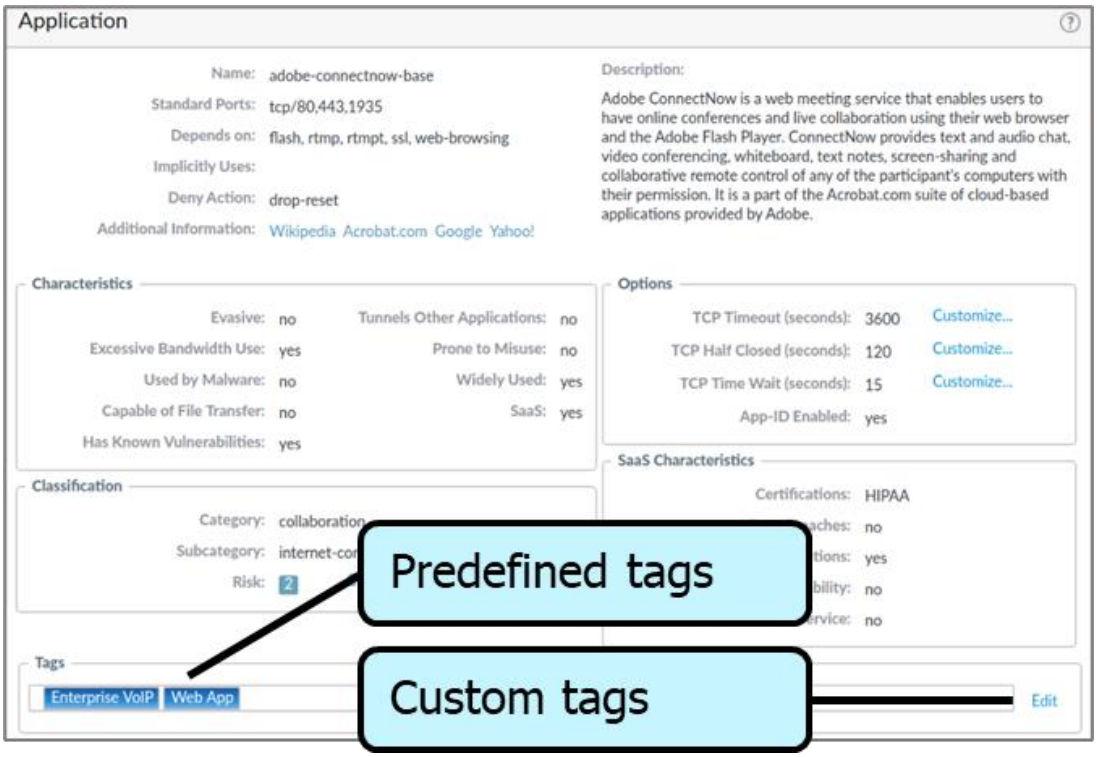

You can use application tags as policy rule match criteria. First, create an application filter using one or more application tags as filter criteria. Then add the application filter to a policy rule. If the tags associated with applications are updated, then the behavior of application filters and policy rules also will be automatically updated. An example Security policy rule with an application filter is shown in the following screenshot:

Application Dependencies

PAN-OS security controls are application-based rather than port-based and protocol-based. For example, rather than permit HTTP port 80 traffic to traverse the firewall, you permit the *web-browsing* application instead. The use of named applications to control network traffic rather than protocol and port number combinations reduces your attack surface.

Because you must work with PAN-OS policies that are application based, you must be aware of concepts of implicit and explicit application dependencies. Some applications depend on other applications. For example, the facebook-base application depends on the web-browsing and ssl applications. Before you can use Facebook in your environment, your firewall policy also must permit the ssl, web-browsing, and

facebook-base applications. The applications that you must explicitly add to your policy rules depend on whether you are working with implicit application dependencies or explicit application dependencies.

The facebook-base application implicitly allows the ssl and web-browsing applications. This “implicit allow” means that you don’t have to explicitly add a policy rule that permits ssl and web-browsing. If you allow the facebook-base application, then the PAN-OS software will implicitly allow ssl and web-browsing too. In the web interface, browse to **Objects > Applications** and then open an application’s details window to view any implicitly allowed applications. Notice the **Implicitly Uses** field in the following screenshot:

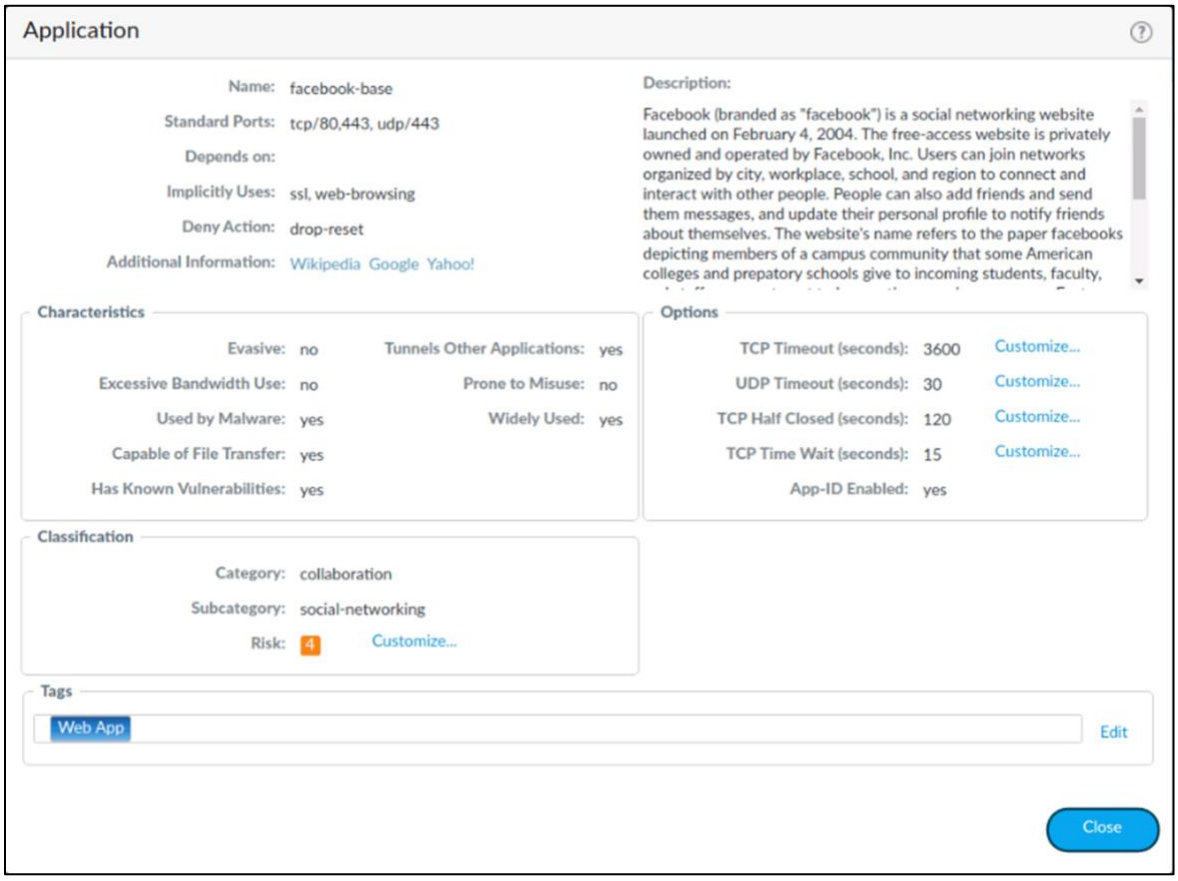

Other dependent applications require that you explicitly add any other required applications to your policy rules. Such applications are said to have explicit application dependencies. For example, to permit the use of Hotmail in your environment, you must permit the hotmail, office365-consumer-access, silverlight, ssl, and web-browsing applications. However, the office365-consumer-access, silverlight, ssl, and web-browsing applications are not implicitly allowed by the hotmail application, which means that you must explicitly add them to your policy rules before your users would be able to use Hotmail. In the web interface, browse to **Objects > Applications** and then open an application's details window to view any explicit application dependencies. Notice the **Depends on** field in the following screenshot:

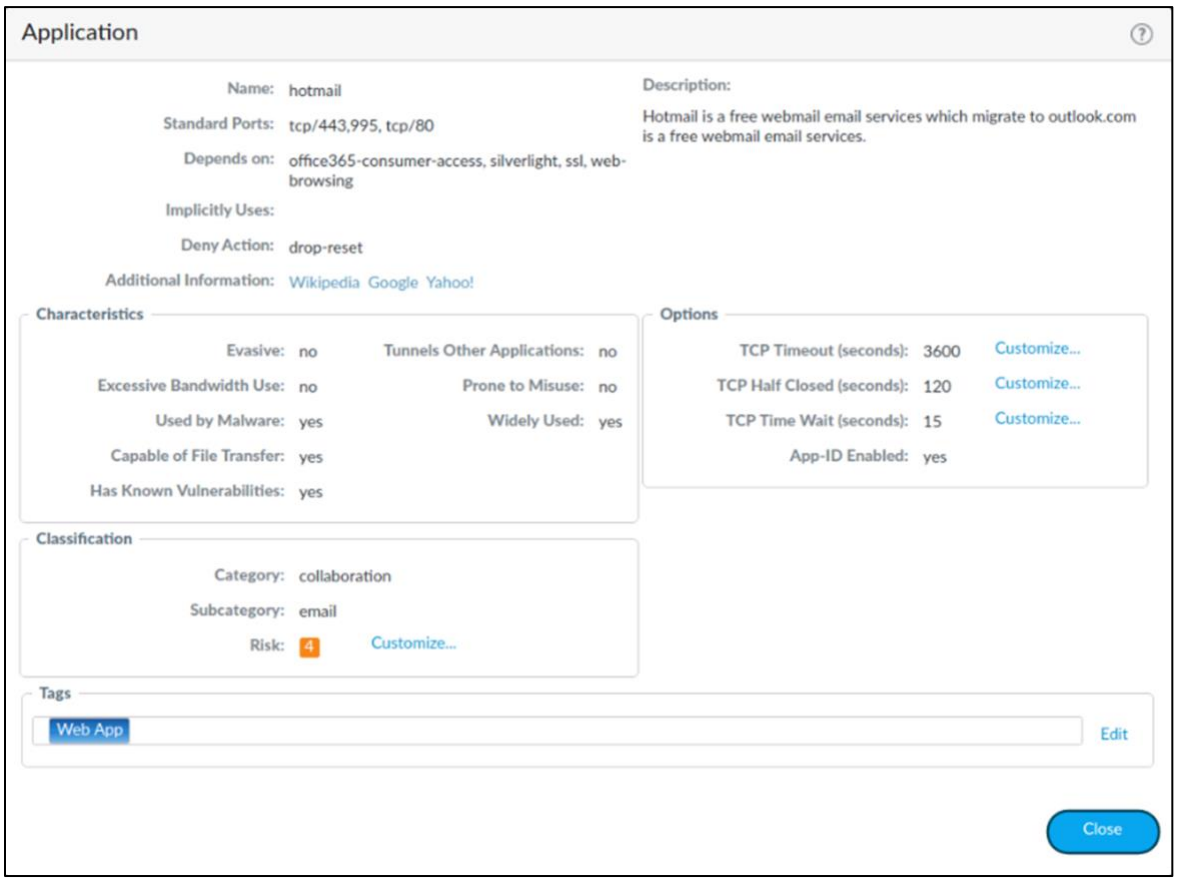

Explicit Application Dependency Resolution Workflows

PAN-OS 9.1 features three different workflows that enable you to configure your policy rules to include any explicit application dependencies. Some workflows are more efficient and streamlined than others.

The most efficient workflow is to determine and resolve explicit application dependencies while you configure a new policy rule. This method is available starting with PAN-OS 9.1. For example, in the web interface, browse to **Policies > Security** and then click **Add** to create a new rule. As you are adding applications on the **Applications** tab, the **Depends on** subpanel lists any explicit application dependencies. You can add the listed applications to the current rule, or to another existing rule, in your policy. When you are finished, commit your configuration. The following screenshot is an example of the new **Depends on** subpanel:

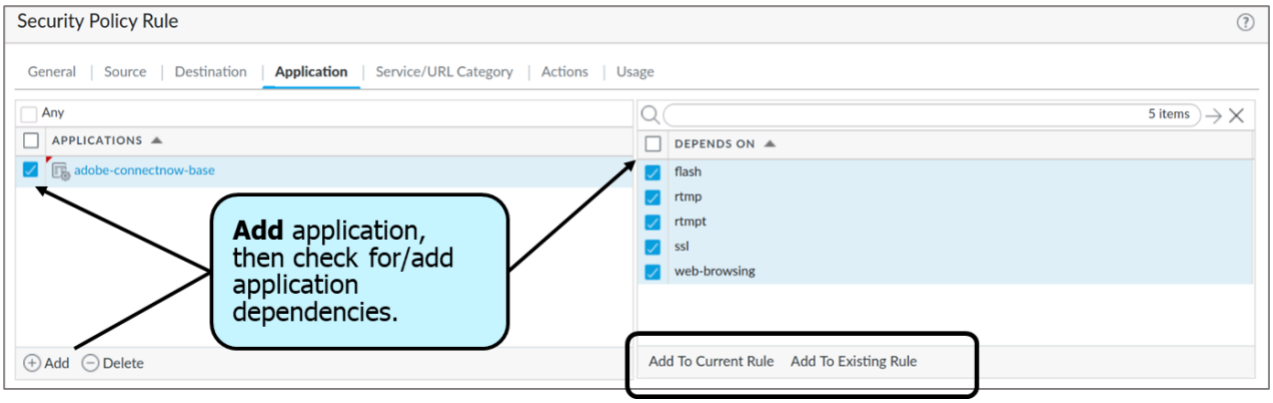

You also can determine explicit application dependencies by updating your policy rules and then performing a commit. This workflow is not as efficient as the previous workflow. When you perform the commit in the web interface, the firewall finishes by presenting you with a **Commit Status** window. Starting with PAN-OS 9.1, the **Commit Status** window includes a new **App Dependency** tab if there are any applications in your policy rules with missing explicit application dependencies. In the example that follows, the "Internet" policy rule permits the adobe-connectnow-base application but does not include the required application dependencies. The **Commit Status** window reports these missing applications. In the **Commit Status** window, you can click **Users_Net-Internet**, edit the rule to add the missing applications, and then commit again:

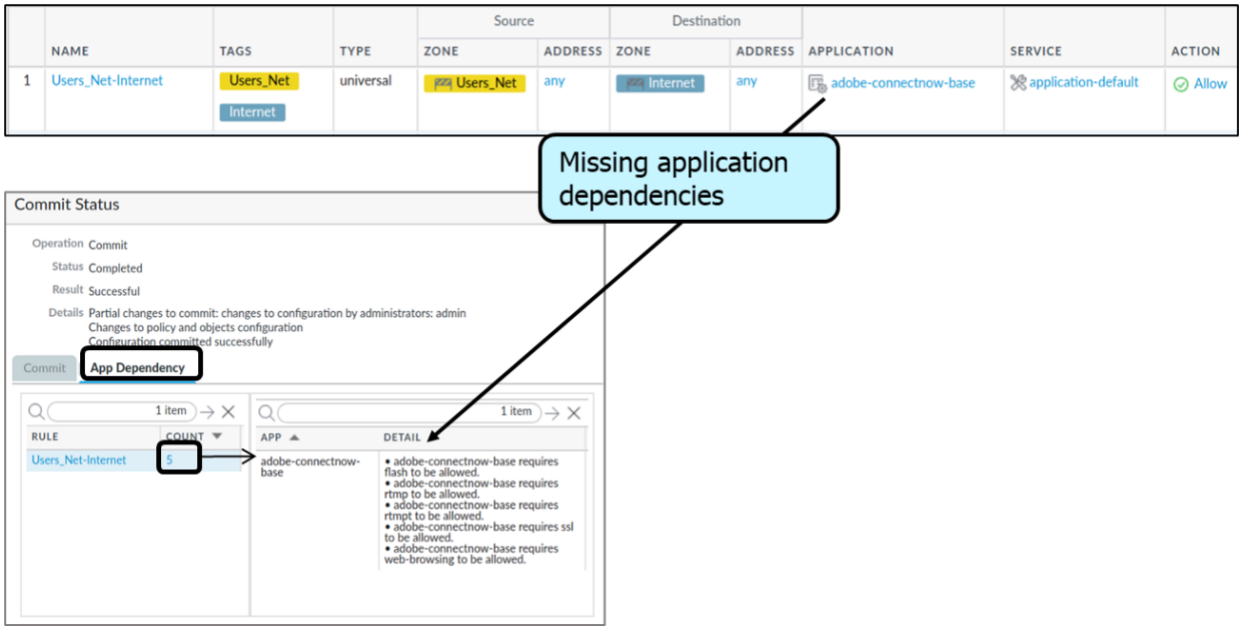

The final workflow is to research the App-ID information presented at **Objects > Applications** before you create or update a policy rule. For example, if you wanted to permit the use of Hotmail, then your App-ID research should tell you that you must add not only the hotmail application to your policy rule but also the office365-consumer-access, silverlight, ssl, and web-browsing applications. Notice the **Depends on** information for Hotmail in the following screenshot:

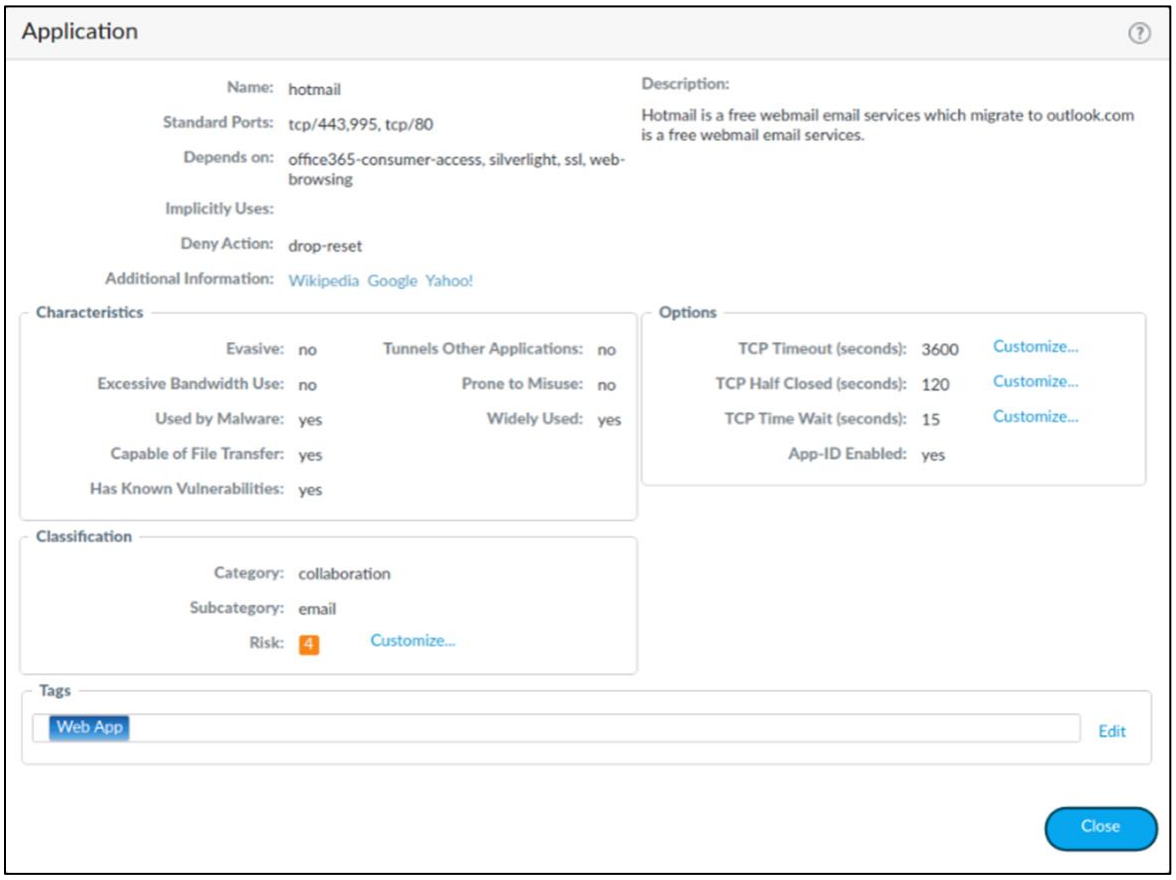

Sample Question

Q4. Which two protocols are implicitly allowed when you select the facebook-base application? (Choose two.)

a) web-browsing

b) chat

c) gaming

d) ssl

Domain 5 – Securing Traffic

5.1 Identify and apply the appropriate Security Profile

Security Profiles

Security Profiles are added to the end of Security policy rules. After a packet has been allowed by the Security policy, Security Profiles are used to scan packets for threats, vulnerabilities, viruses, spyware, malicious URLs, data exfiltration, and exploitation software. Traffic also can be scanned for suspicious file uploads.

A Security Profile Group can be created that includes one or more Security Profiles, which simplifies the task of adding Security Profiles to a Security policy rule.

The following table describes the Security Profile types:

Type	Description
Antivirus	Detects infected files being transferred within the application or protocol
Anti-Spyware	Detects spyware downloads and command-and-control traffic from previously installed spyware
Vulnerability Protection	Detects attempts to exploit known software vulnerabilities
URL Filtering	Classifies and controls web browsing based on website content
File Blocking	Tracks and blocks file uploads and downloads based on file type and application
Data Filtering	Identifies and blocks transfer of specific data patterns found in network traffic **Note:** This type will not be discussed further in this section.
WildFire Analysis	Forwards unknown files and URL links to the WildFire® service for malware analysis. **Note:** This type will not be discussed further in this section.

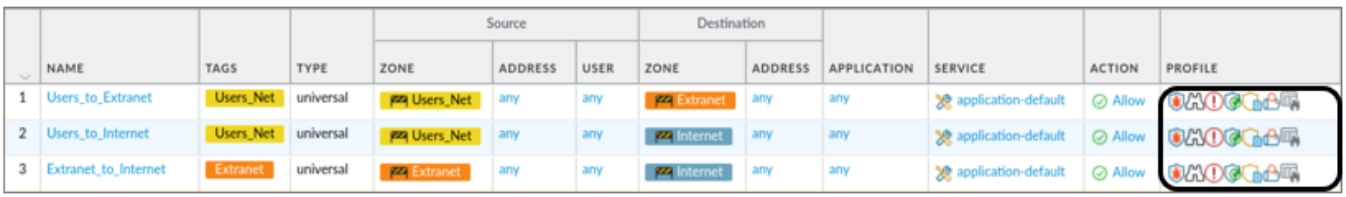

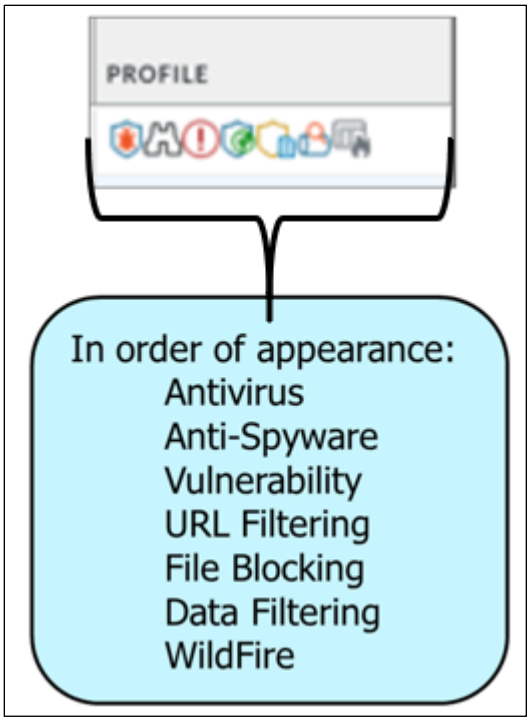

Threat Logs

Threats are recorded and logged in the Threat log. A Threat log display entries when traffic matches one of the Security Profiles attached to a Security policy 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. The Threat log is used as the source of information that is displayed on the **ACC** tab (Application Control Center).

Threat levels are based on severity. There are five levels of severity:

- **Critical:** Critical threats are serious threats such as those that affect default installations of widely deployed software and result in the compromise of servers. Critical threats include those where 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**: High threats are those that can become critical but have mitigating factors; for example, they might be difficult to exploit, do not result in elevated privileges, or do not have a large victim pool.

- **Medium**: Medium threats are minor threats and those that pose minimal impact. Examples include DoS attacks that do not compromise the target or exploits that require an attacker to reside on the same LAN as the victim. Medium threats affect only non-standard configurations or obscure applications, or provide very limited access.

- **Low**: Low threats are warning-level threats that have little impact on an organization's infrastructure. They usually require local or physical system access and often might result in victim privacy or DoS issues and information leakage. Data Filtering Profile matches are logged as Low.
- **Informational**: Informational threats are suspicious events that do not pose an immediate threat but that are reported to call attention to deeper problems that could exist. URL Filtering log entries are logged as Informational. Log entries with any verdict and an action set to block also are logged as Informational.

Antivirus Security Profiles

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

The *default* profile inspects all listed protocol decoders for viruses and generates alerts for SMTP, IMAP, and POP3 protocols while blocking 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.

Customized profiles can be used to minimize antivirus inspection for traffic between more trusted security zones. They also can be used to maximize the inspection of traffic received from less untrusted zones, such as the internet, and 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. Signatures are quickly created as threats are discovered by WildFire and then are integrated into the standard antivirus signatures that can be downloaded daily by Threat Prevention subscribers (sub-hourly for WildFire subscribers).

Anti-Spyware Security Profiles

Anti-Spyware Security Profiles block spyware on compromised hosts from trying to communicate with external command-and-control (C2) servers, thus enabling you to detect malicious traffic leaving the network from infected clients. You can apply various levels of protection between security zones. For example, you might want to have custom Anti-Spyware Profiles that minimize inspection between more trusted zones while maximizing inspection on traffic received from less trusted zone such as internet-facing zones.

Vulnerability Protection Security Profiles

Vulnerability Protection Security Profiles stop attempts to exploit system flaws or gain unauthorized access to systems. Anti-Spyware Security Profiles identify infected hosts as traffic leaves the network, but Vulnerability Protection Security Profiles protect against threats entering the network. For example, Vulnerability Protection Security Profiles protect against buffer overflows, illegal code execution, and other attempts to exploit system vulnerabilities. The *default* Vulnerability Protection Security Profile protects clients and servers from all known critical-, high-, and medium-severity threats. You also can create exceptions that enable you to change the response to a specific signature.

URL Filtering Security Profiles

The URL Filtering Security Profile determines 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 Security Profile. All allowed traffic will not be logged by default. You can customize the URL Filtering Security Profile with custom site access settings for each category, or use the predefined *default* URL Filtering Security 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.

For each URL category, select **User Credential Submissions** to allow or disallow users from submitting valid corporate credentials to a URL in that category. This action will prevent credential phishing.

Management of 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 automatically are set to also block user credential submissions.

File Blocking Security Profiles

A Security policy can include specification of a File Blocking Profile that blocks selected file types from being uploaded or downloaded, or generates an alert when the specified file types are detected.

Sample Question

Q1. What are two benefits of Vulnerability Protection Security Profiles? (Choose two.)

a) prevent compromised hosts from trying to communicate with external C2 servers

b) protect against viruses, worms, and Trojans

c) prevent exploitation of system flaws

d) prevent unauthorized access to systems

Advanced Persistent Threats

Threats to your organization are growing in complexity and capability. Advanced persistent threats (APTs) represent the most difficult challenge to the security professional.

Security Policy and Profiles

The primary firewall tools protecting users from threats are Security policy rules combined with Security Profiles implementing specific protections.

The first steps in creating a Security policy are found here:
https://docs.paloaltonetworks.com/pan-os/10-0/pan-os-admin/getting-started/set-up-a-basic-security-policy.html

The completion of these steps provides only a basic setup that is not comprehensive enough to protect your network. The next phase is here:
https://docs.paloaltonetworks.com/best-practices/10-0/internet-gateway-best-practices.html

Security Profiles are an important aspect of protection, detection, and prevention for specific types of threats. See the document at the following link for more details:
https://docs.paloaltonetworks.com/best-practices/10-0/internet-gateway-best-practices/best-practice-internet-gateway-security-policy/create-best-practice-security-profiles.html

Sample Questions

Q2. A URL Filtering Profile is part of which type of identification?

a) App-ID

b) Content-ID

c) User-ID

d) Service

Q3. Which stage of the attack lifecycle is most likely to be stopped by dividing the network into separate security zones?

a) Reconnaissance

b) Execution

c) Lateral movement

d) Data exfiltration

Q4. Which component can tell you if an attack is an APT or a broad attack designed to produce a botnet for future abuse?

a) next-generation firewall

b) WildFire

c) MineMeld

d) AutoFocus

Denial-of-Service Attack (DoS)

A denial-of-service (DoS) attack attempts to make a network device or resource unavailable to legitimate users by disrupting services. These attacks usually come from the Internet but can come from misconfigured or compromised internal devices. The typical method is to flood the target with resource requests until the requests consume all of the target's available resources—memory, CPU cycles, and bandwidth—and the target becomes unavailable. Typical targets are Internet-facing devices users can access from outside the corporate network, such as web servers and database servers. As part of a layered approach to DoS protection, Palo Alto Networks firewalls provide three DoS attack mitigation tools:

1. Zone Protection Profiles

Zone Protection Profiles: Protect only new sessions in ingress zones to defend against flood attacks by limiting the connections per second (CPS) to the firewall. The profiles also provide protection against reconnaissance (port scans and host sweeps), packet-based attacks, and Layer 2 protocol-based attacks. See the Zone Protection Profiles section in 1.15 Given a scenario, identify ways to mitigate resource exhaustion (because of denial-of-service) in application servers for a complete discussion.

2. Dos Protection Profiles and Policy Rules

DoS Protection Profiles and policy rules: Provide granular protection of specific, critical devices for new sessions. Classified policies protect individual devices by limiting the CPS for a specific device. Aggregate policies limit the total CPS for a group of devices, but don't limit the CPS for a particular device in that group to less than the total allowed for the group, which means that one device still may receive the majority of the connection requests. See the *DoS Protection Profile* section in *1.15 Given a scenario, identify ways to mitigate resource exhaustion (because of denial-of-service) in application servers* for a complete discussion.

3. Packet Buffer Protection

Packet Buffer Protection: Defends 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 in a Zone Protection Profile or DoS Protection Profile or policy rule, Packet Buffer Protection defends ingress zones. Although zone and DoS protection apply to new sessions (connections) and are granular, Packet Buffer Protection applies to existing sessions and is global.

You configure Packet Buffer Protection globally to protect the entire firewall and each zone:

- Global Packet Buffer Protection: Monitors sessions from all zones (regardless of whether Packet Buffer Protection is enabled in a zone) to see how those sessions use the packet buffer. You must configure Packet Buffer Protection globally (**Device > Setup > Session Settings**) to protect the firewall and to enable it on individual zones. When packet buffer consumption reaches the configured Activate percentage, the firewall used Random Early Drop (RED) to drop packets from the offending sessions (the firewall doesn't drop complete sessions at the global level).

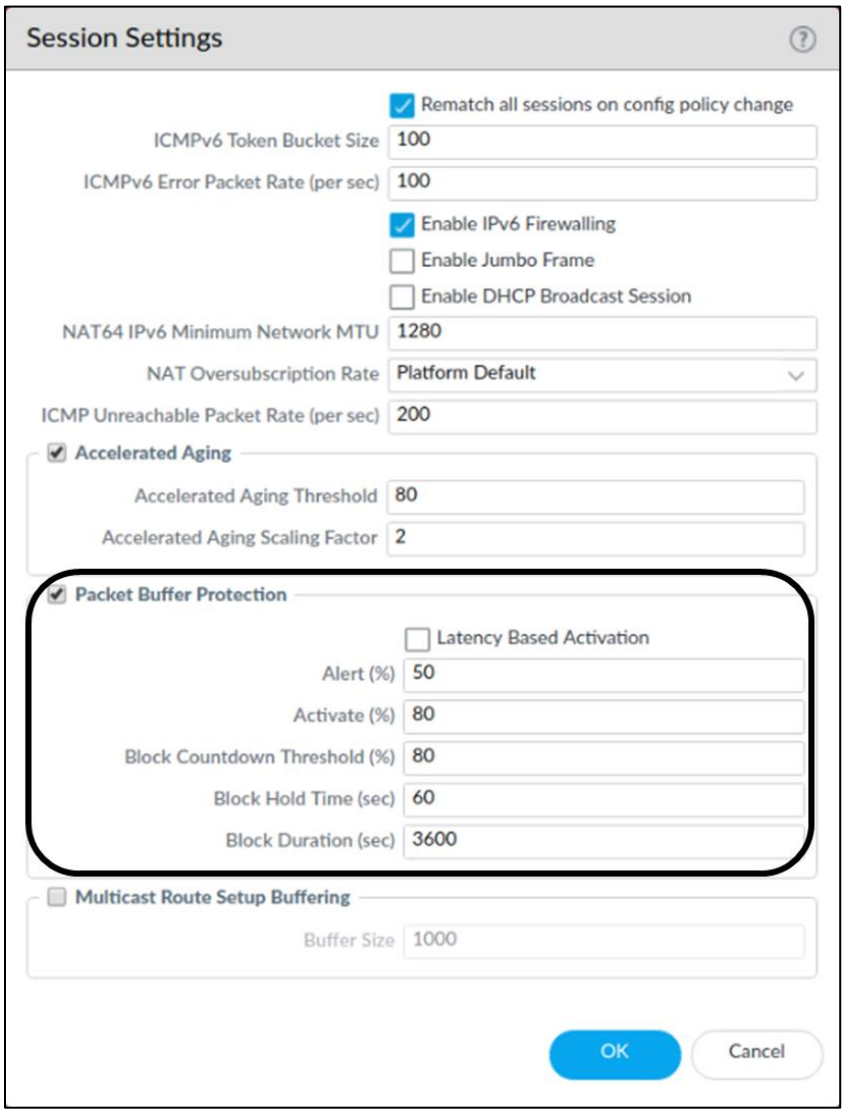

- Per-Zone Packet Buffer Protection: Provides a second level of protection for each zone (**Network > Zones**). When packet buffer consumption crosses the Activate threshold and global protection begins to apply RED to session traffic it starts the Block Hold Time timer. The Block Hold Time is the amount of time, in seconds, that the offending session can continue before the firewall blocks the entire session. The offending session remains blocked until the Block Duration time expires.

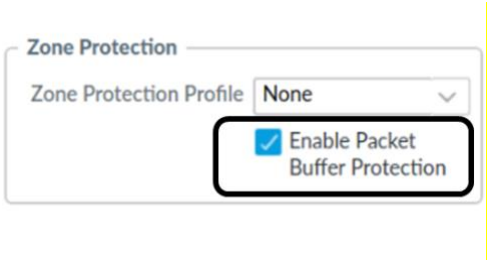

Sample Questions

Q5. Packet Buffer Protection defends against which type of denial-of-service attack?

a) from distributed sessions

b) from a single App-ID source

c) from multiple App-ID sources

d) from a single session

Q6. Which defense is turned on when a Packet Buffer Protection event is detected? (Choose two.)

a) SYN cookie management of attacking session traffic

b) Global Random Early Drop of packets from the attacking session

c) block all packets from the attacking session for the configured duration if the attack persists for a certain configured time

d) block all packets from the attacking IP address for the configured duration if the attack persists for a certain configured time

Denial-of-Service Protection

PAN-OS software provides protection not only through use of Security policy rules and Security Profiles, which use signatures and heuristics to identify attacks. PAN-OS software also provides denial-of-service (DoS) protection, which is based on analysis of packet headers to detect threats rather than signatures.

A DoS attack attempts to make network devices unreachable by disrupting services. These attacks usually come from the internet but can come from misconfigured or compromised internal devices. The typical method is to flood the target with resource requests until the requests consume all the target's available resources: memory, CPU cycles, and bandwidth. Typical targets are internet-facing devices that users can access from outside the corporate network such as web servers and database servers. Palo Alto Networks firewalls provide the following three DoS attack mitigation tools as part of a layered approach to DoS protection. Note that packet buffer protection will not be described in this section.

- **Zone Protection Profiles**: Apply only to new sessions in ingress zones and provide broad protection against flood attacks by limiting the connections-per-second (CPS) to the firewall, plus protection

against reconnaissance (port scans and host sweeps), packet-based attacks, and Layer 2 protocol-based attacks.

- **DoS Protection Profiles and policy rules**: Provide granular protection of specific, critical devices for new sessions. Classified profiles protect individual devices by limiting the CPS for a specific device or specific devices. Aggregate profiles limit the total CPS for a group of devices but don't limit the CPS for a particular device in the group to less than the total allowed for the group, so one device still might receive most of the connection requests.
- **Packet buffer protection:** Protects against single-session DoS attacks that attempt to overwhelm the firewall's packet buffer

Zone Protection Profiles

A Zone Protection Profile is applied to an ingress zone. It offers protection against floods, reconnaissance attacks, and other packet-based attacks. Zone protection is broad-based protection and is not designed to protect a specific end host or traffic going to a particular destination zone. Only a single Zone Protection Profile can be applied to a zone. Zone protection is enforced only when there is no session match for the packet because zone protection is based on new CPS, not on packets per second (pps). If the packet matches an existing session, it will bypass the Zone Protection Profiles.

Flood Attack Protection

Zone Protection Profiles protect against of five types of floods:

- SYN (TCP)
- UDP
- ICMP
- ICMPv6
- Other IP

Flood Protection Activate Rates

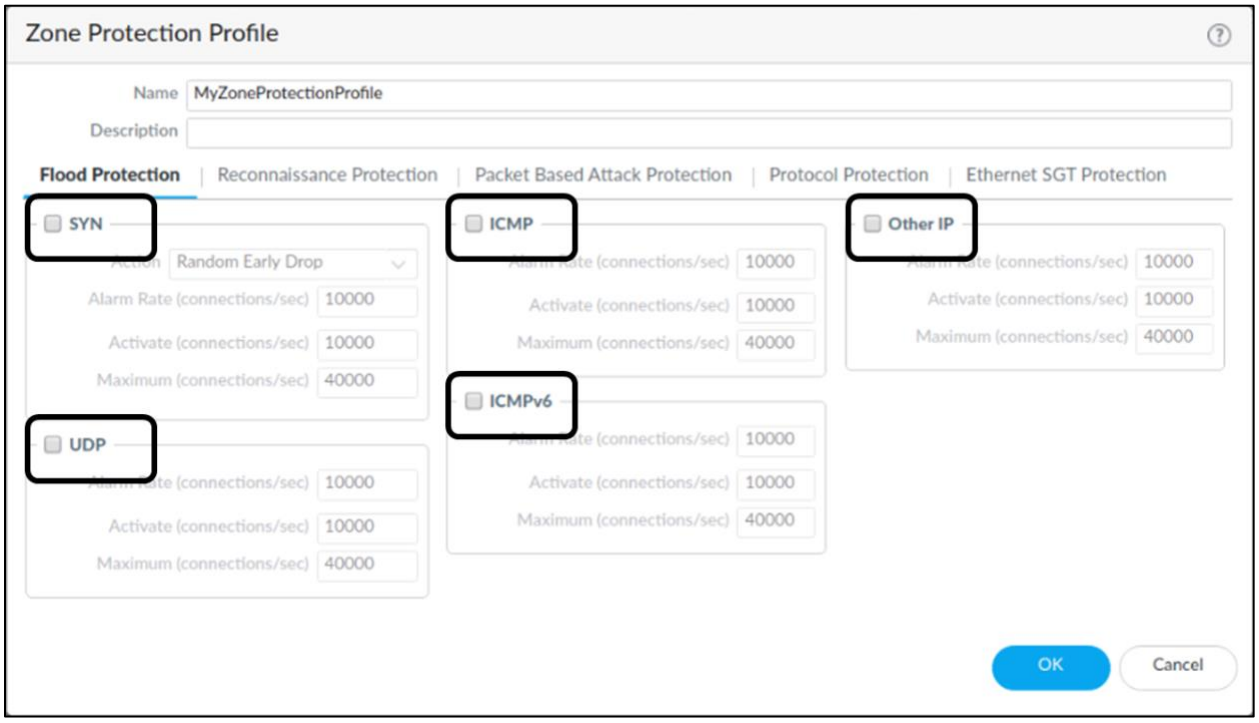

SYN Random Early Drop

This feature causes TCP SYN packets to be dropped to mitigate a flood attack. When the flow exceeds the **Activate** rate threshold, the firewall drops individual SYN packets randomly to restrict the flow. When the flow exceeds the **Maximum** rate threshold, 100% of incoming SYN packets are dropped.

SYN Cookies

This feature causes the firewall to act like a proxy, intercept the TCP SYN, generate a cookie on behalf of the server to which the SYN was directed, and send a SYN-ACK with the cookie to the original source. Only when the source returns an ACK with the cookie to the firewall does the firewall consider the source valid and forward the SYN to the server. This is the preferred configuration option.

UDP

UDP flood protection is activated when the number of UDP packets (not matching an existing session) the zone receives per second exceeds the **Activate** threshold. The firewall uses an algorithm to progressively drop more packets as the attack rate increases, until the rate reaches the **Maximum** rate. The firewall stops dropping the UDP packets if the incoming rate drops below the **Activate** threshold.

ICMP

ICMP flood protection is activated when the number of ICMP packets (not matching an existing session) the zone receives per second exceeds the **Activate** threshold. The firewall uses an algorithm to progressively drop more packets as the attack rate increases, until the rate reaches the **Maximum** rate. The firewall stops dropping the ICMP packets if the incoming rate drops below the **Activate** threshold.

ICMPv6

ICMPv6 flood protection is activated when the number of ICMPv6 packets (not matching an existing session) the zone receives per second exceeds the **Activate** threshold. The firewall uses an algorithm to progressively drop more packets as the attack rate increases, until the rate reaches the **Maximum** rate. The firewall stops dropping the ICMPv6 packets if the incoming rate drops below the **Activate** threshold.

Other IP

Other IP flood protection is activated when the number of non-IP packets (not matching an existing session) the zone receives per second exceeds the **Activate** threshold. The firewall uses an algorithm to progressively drop more packets as the attack rate increases, until the rate reaches the **Maximum** rate. The firewall stops dropping the Other IP packets if the incoming rate drops below the **Activate** threshold.

Reconnaissance Attack Protection

Reconnaissance protection protects against reconnaissance attacks, which are the first type of attacks within a cyberattack lifecycle. During the first stage of the attack lifecycle, cyberattackers carefully plan their method of attack. They research, identify, and select targets within an organization such as human resources and financial personnel that will enable them to meet their objectives. Attackers can gather intelligence through publicly available sources such as Twitter, LinkedIn, and corporate websites , all the places where a company will share information about itself. The cyberattackers also will scan for vulnerabilities that can be exploited within the target network (services and applications), and map out resources that they can take advantage of.

Reconnaissance attacks are prevented by:

- Performing continuous inspection of network traffic flows to detect and prevent port scans and host sweeps
- Implementing security awareness by limiting what should be posted on the internet: Examples of content that should not be posted are sensitive documents, customer lists, event attendees, job roles, and responsibilities

See Palo Alto Networks documentation for more detail about these complicated types of attacks.

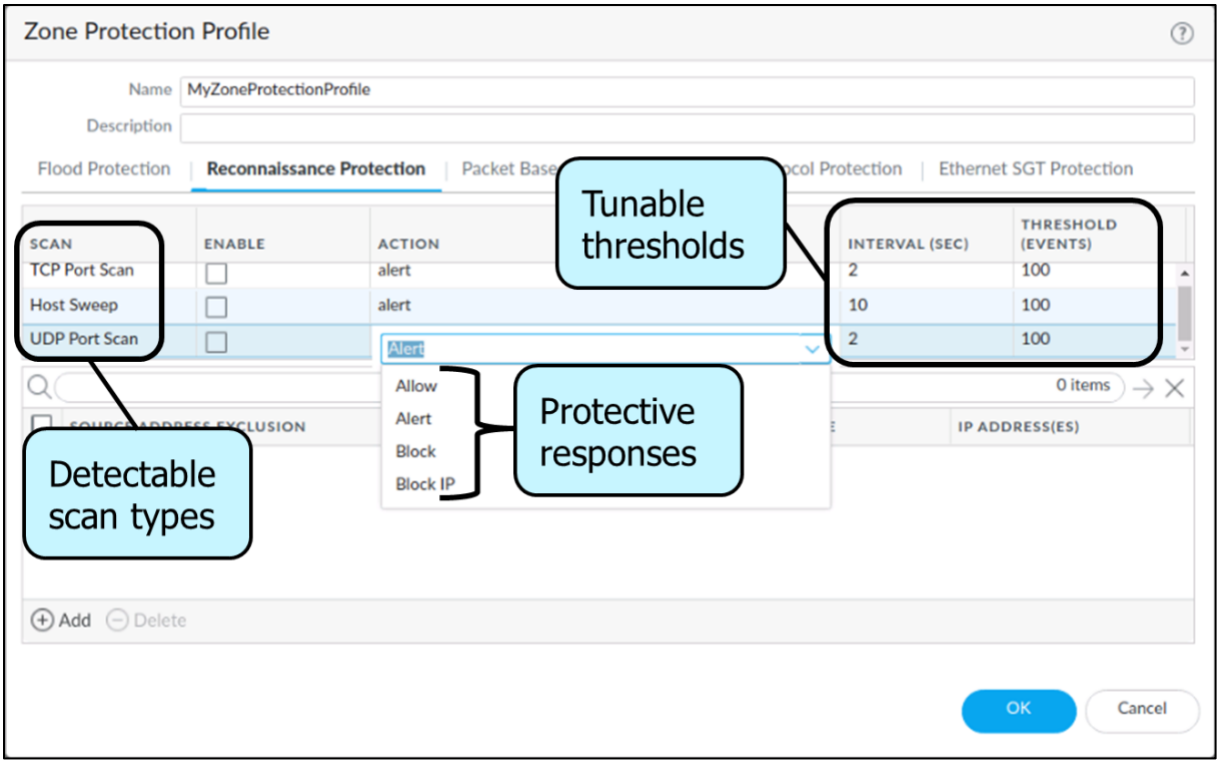

Packet-Based Attack Protection

There are many types of packet-based attack protection. Each one will not be covered in detail in this section. Please refer to Palo Alto Networks documentation to obtain more detail about these complicated types of attacks.

The five major categories of packet-based attack protection are:

- IP Drop
- TCP Drop
- ICMP Drop
- IPv6 Drop
- ICMPv6 Drop

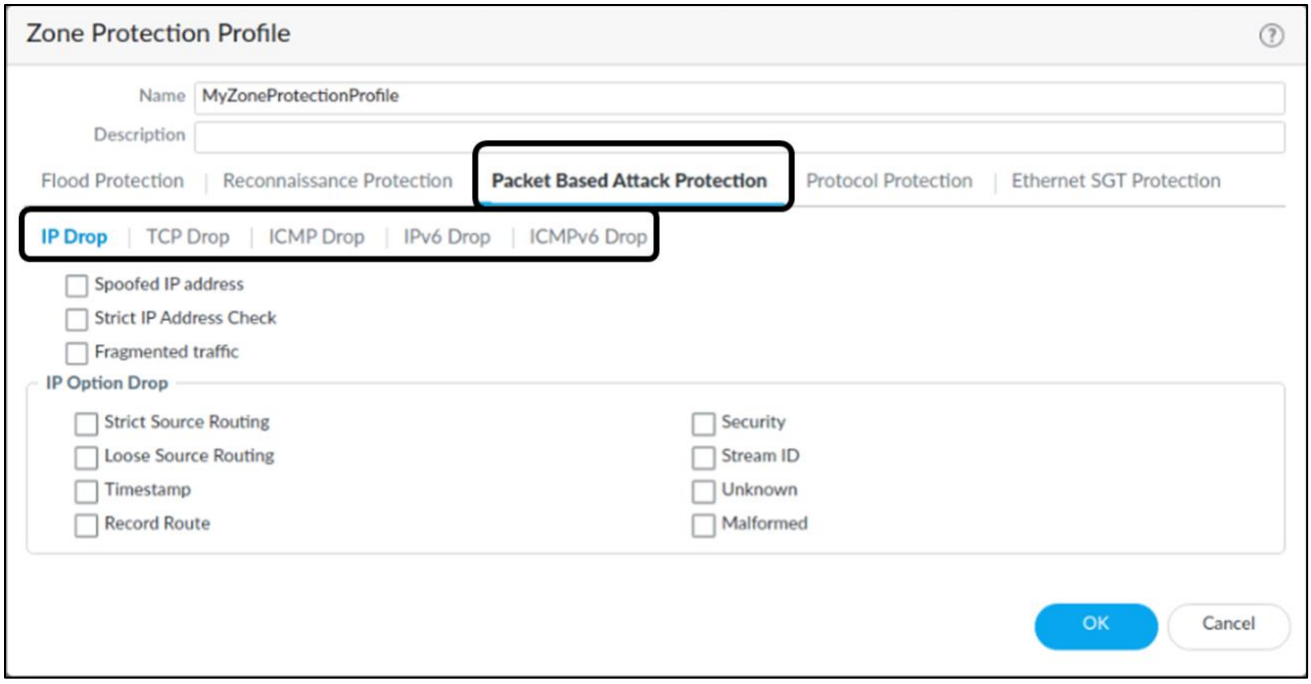

Packet-based attacks take many forms. Zone Protection Profiles check IP, TCP, ICMP, IPv6, and ICMPv6 packet headers and protect a zone by:

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

Select the drop characteristics for each packet type when you configure packet-based attack protection. The best practices for each IP protocol are:

- **IP Drop**: Drop **Unknown** and **Malformed** packets. Also drop **Strict Source Routing** and **Loose Source Routing** because allowing these options permits adversaries to bypass Security policy rules that use the destination IP address as the matching criterion. For internal zones only, check **Spoofed IP Address** so only traffic with a source address that matches the firewall routing table can access the zone.
- **TCP Drop**: Retain the default **TCP SYN with Data** and **TCP SYNACK with Data** drops, drop **Mismatched overlapping TCP segment** and **Split Handshake** packets, and strip the **TCP Timestamp** from packets.

Enable the **Rematch Sessions** option as a best practice. It applies committed newly configured or edited Security policy rules to existing sessions. However, if you configure **Tunnel Content Inspection** on a zone and **Rematch Sessions** is enabled, you also must disable **Reject Non-SYN TCP** (change the selection from **Global** to **No**). If you don't disable that option, when you enable or edit a **Tunnel Content Inspection** policy the firewall will drop all existing tunnel sessions. Create a separate Zone Protection Profile to disable **Reject Non-SYN TCP** only on zones that have **Tunnel Content Inspection** policies and only when you enable **Rematch Sessions**.

- **ICMP Drop**: There are no standard best practice settings because dropping of ICMP packets depends on how you use ICMP (or if you use ICMP). For example, if you want to block ping activity, you can block **ICMP Ping ID 0**.
- **IPv6 Drop**: If compliance matters, ensure that the firewall drops packets with non-compliant routing headers, extensions, etc.
- **ICMPv6 Drop**: If compliance matters, ensure that the firewall drops certain packets if the packets don't match a Security policy rule.

Protocol Attack Protection

In a Zone Protection Profile, **Protocol Protection** defends against non-IP protocol-based attacks. Enable **Protocol Protection** to block or allow non-IP protocols between security zones on a Layer 2 VLAN or on a virtual wire, or between interfaces within a single zone on a Layer 2 VLAN (Layer 3 interfaces and zones drop non-IP protocols, so non-IP **Protocol Protection** doesn't apply).

Configure **Protocol Protection** to reduce security risks and facilitate regulatory compliance by preventing less secure protocols from entering a zone or an interface in a zone. If you don't configure a Zone Protection Profile that prevents non-IP protocols in the same zone from going from one Layer 2 interface to another, the firewall allows the traffic because of the default intrazone allow Security policy rule. You can create a Zone Protection Profile that blocks protocols such as LLDP within a zone to prevent discovery of networks reachable through other zone interfaces.

If you need to discover which non-IP protocols are running on your network, use monitoring tools such as NetFlow, Wireshark, or other third-party tools. Examples of non-IP protocols you can block or allow are LLDP, NetBEUI, Spanning Tree, and Supervisory Control and Data Acquisition (SCADA) systems such as Generic Object Oriented Substation Event (GOOSE), among many others.

Create an **Exclude List** or an **Include List** to configure **Protocol Protection** for a zone. The **Exclude List** is a blacklist: The firewall blocks all the protocols you place in the **Exclude List** and allows all other protocols. The **Include List** is a whitelist: The firewall allows only the protocols you specify in the list and blocks all other protocols. Use include lists for **Protocol Protection** instead of exclude lists. Include lists specifically sanction only the protocols you want to allow and block the protocols you don't need or didn't know were on your network, which reduces the attack surface and blocks unknown traffic. A list supports up to 64 Ethertype entries, each identified by its IEEE hexadecimal Ethertype code. 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 member because AE interface members are treated as a group.

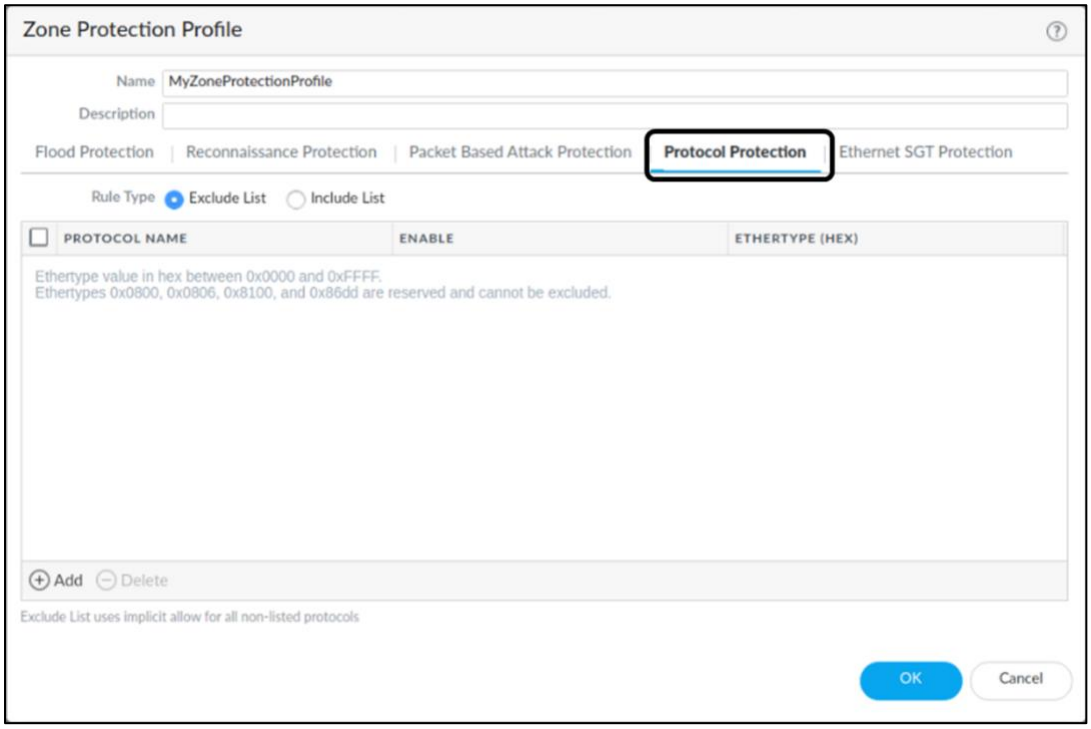

Ethernet SGT Protection

When your firewall is part of a Cisco TrustSec network, the firewall now can inspect headers with 802.1Q (Ethertype 0x8909) for specific Layer 2 Security Group Tag (SGT) values and drop the packet if the SGT matches the list configured in the Zone Protection Profile attached to the interface.

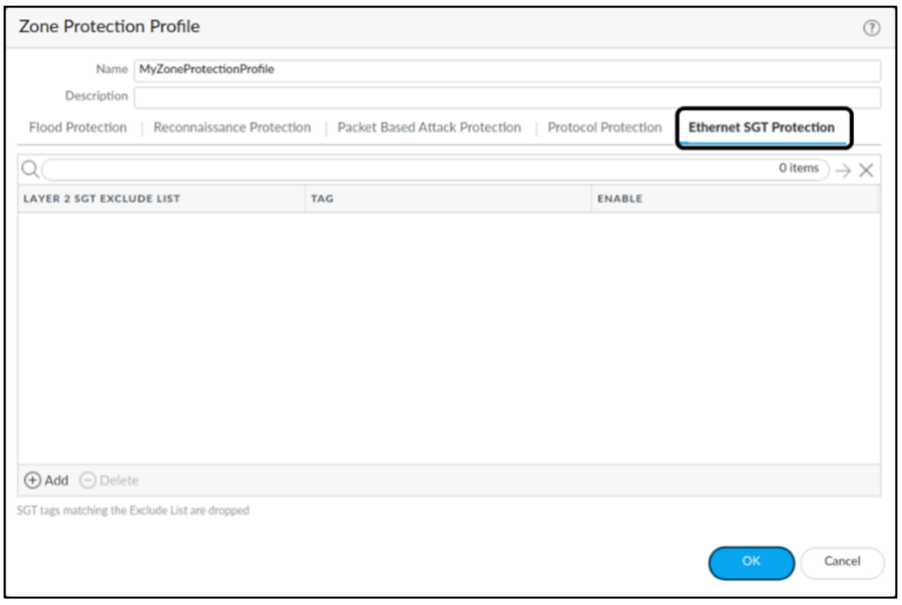

DoS Protection Profiles and Policies

DoS Protection Profiles and DoS Protection policy rules combine to protect specific groups of critical resources and individual critical resources against session floods. Compared to Zone Protection Profiles, which protect entire zones from flood attacks, DoS protection provides granular defense for specific systems, especially critical systems that users access from the internet and often are attack targets, such as web servers and database servers. Apply both types of protection because if you apply only a Zone Protection Profile, then a DoS attack that targets a particular system in the zone can succeed if the total CPS doesn't exceed the zone's **Activate** and **Maximum** rates. DoS protection is resource-intensive, so use it only for critical systems. DoS Protection Profiles specify flood thresholds, similarly to Zone Protection Profiles. DoS Protection policy rules determine the devices, users, zones, and services to which DoS Protection Profiles apply. See Palo Alto Networks documentation for more detail about these complicated types of attacks.

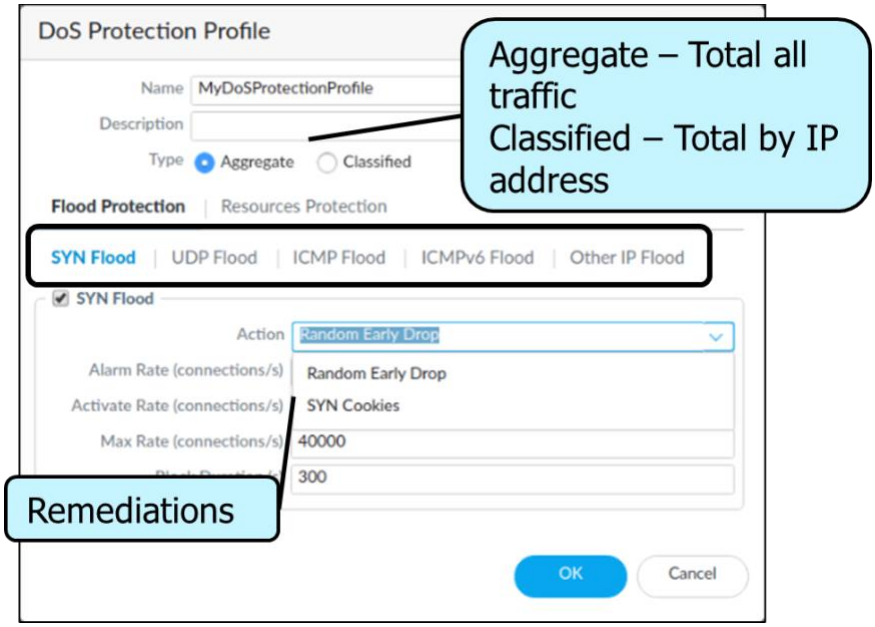

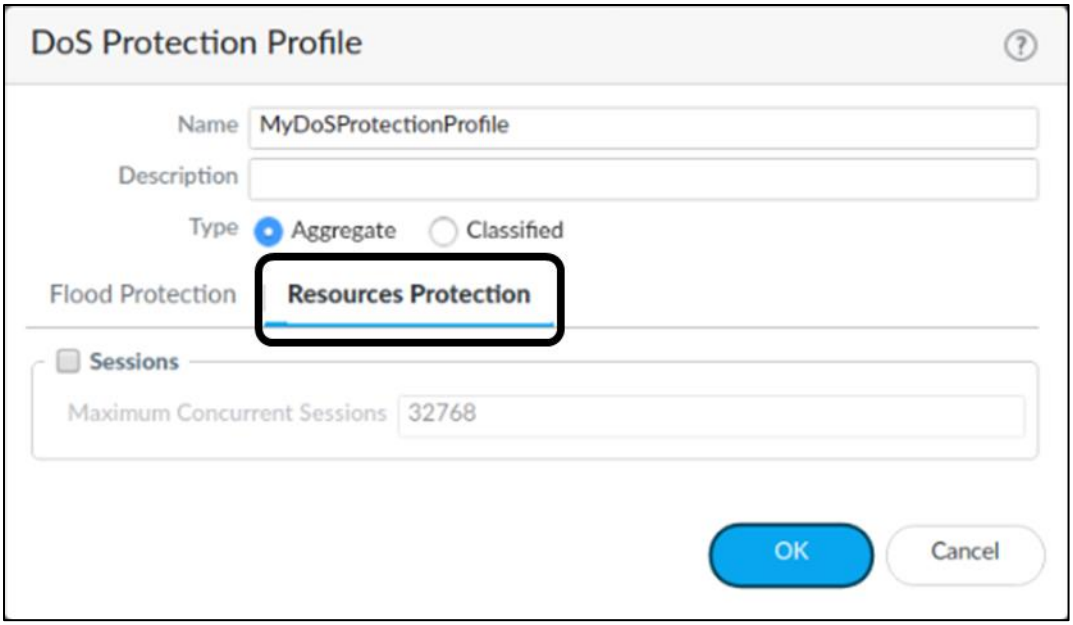

Sample Question

Q7. What are the two components of Denial-of-Service Protection? (Choose two.)

a) Zone Protection Profile

b) DoS Protection Profile and policy rules

c) load protection

d) reconnaissance protection

Q8. Which statement describes the new machine learning capabilities implemented within security profiles introduced in PAN-OS 10.0?

a) Machine learning can be performed by the firewall on the stream of data passing through it, allowing threats to be blocked without signatures.

b) Machine learnt models can be implemented by the firewall on the stream of data passing through it, allowing threats to be blocked without signatures.

c) Machine learnt models can be implemented by the firewall, but only to detect threats after they have passed through the firewall.

d) Machine learning can be performed by the firewall on the stream of data passing through it, identifying threats that have already passed through the firewall.

5.2 Identify the difference between Security policy actions and Security Profile actions

Security Policy Actions and Security Profile Actions

When packets traverse a firewall, they are inspected in two primary stages:

- Security Policy Stage
- Security Profile Stage

In the Security Policy Stage, packets must meet all the criteria in a Security policy rule to match the Security policy rule. If all the criteria match, the Security policy rule's action is applied. If the Security policy action is "allow," the packet is inspected by the Security Profiles attached to the Security policy rule. If all the Security Profile criteria do not match, or the Security policy is any action other than "allow," the packet is evaluated against the next Security policy rule, and so on. You can create a Security Profile Group that includes one or more Security Profiles, which simplifies the task of adding Security Profiles to a Security policy rule.

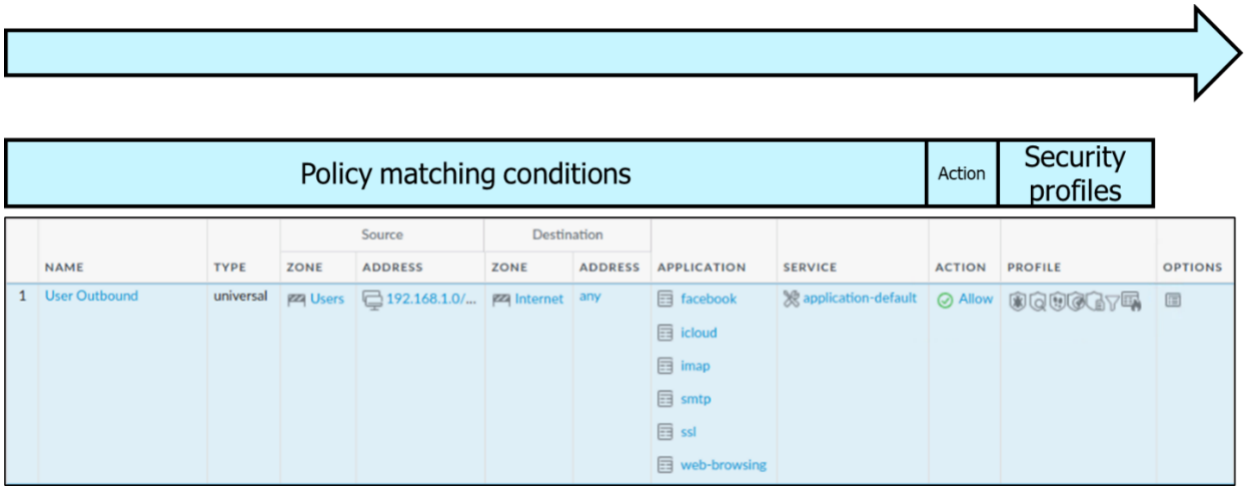

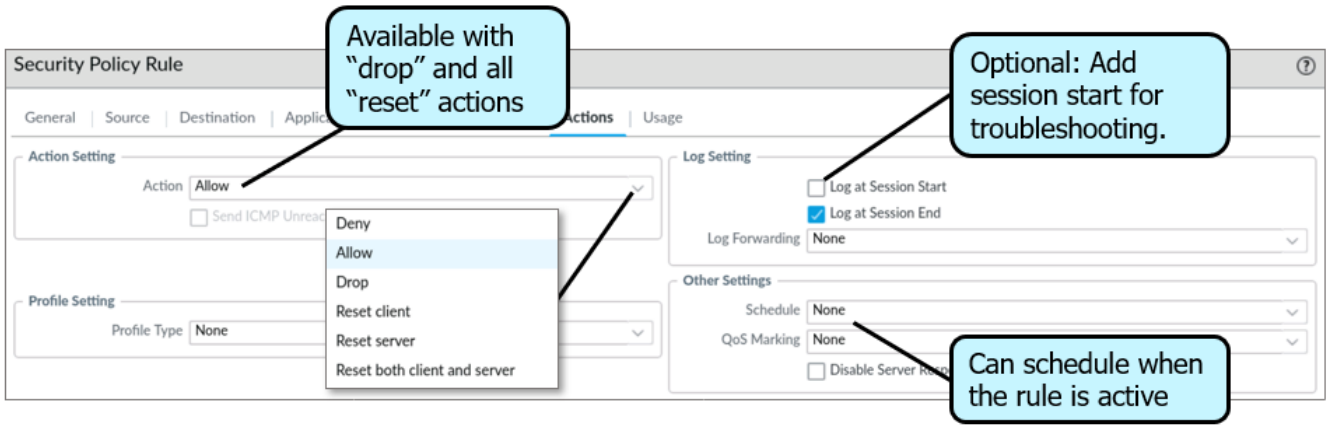

Antivirus Security Profile Actions

The *default* profile inspects the listed protocol decoders for viruses and generates alerts for SMTP, IMAP, and POP3 protocols while blocking 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; see the following table:

Action	Description
Default	For each threat signature and Antivirus signature that is defined by Palo Alto Networks, a default action is specified internally. The default action typically is an "alert" or a "reset-both." The default action is displayed in parentheses, 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 the client and server ends. For UDP, drops the connection.

You also can use customized profiles to minimize antivirus inspection for traffic between trusted security zones. They also can be used to maximize the inspection of traffic received from more untrusted zones such as the internet and of traffic sent to highly sensitive destinations such as server farms. The Palo Alto Networks WildFire product also provides signatures for persistent threats that are more evasive and have not yet been discovered by other antivirus solutions. As threats are discovered by WildFire, signatures are quickly created and then integrated into the standard antivirus signatures that can be downloaded daily by Threat Prevention subscribers (sub-hourly for WildFire subscribers).

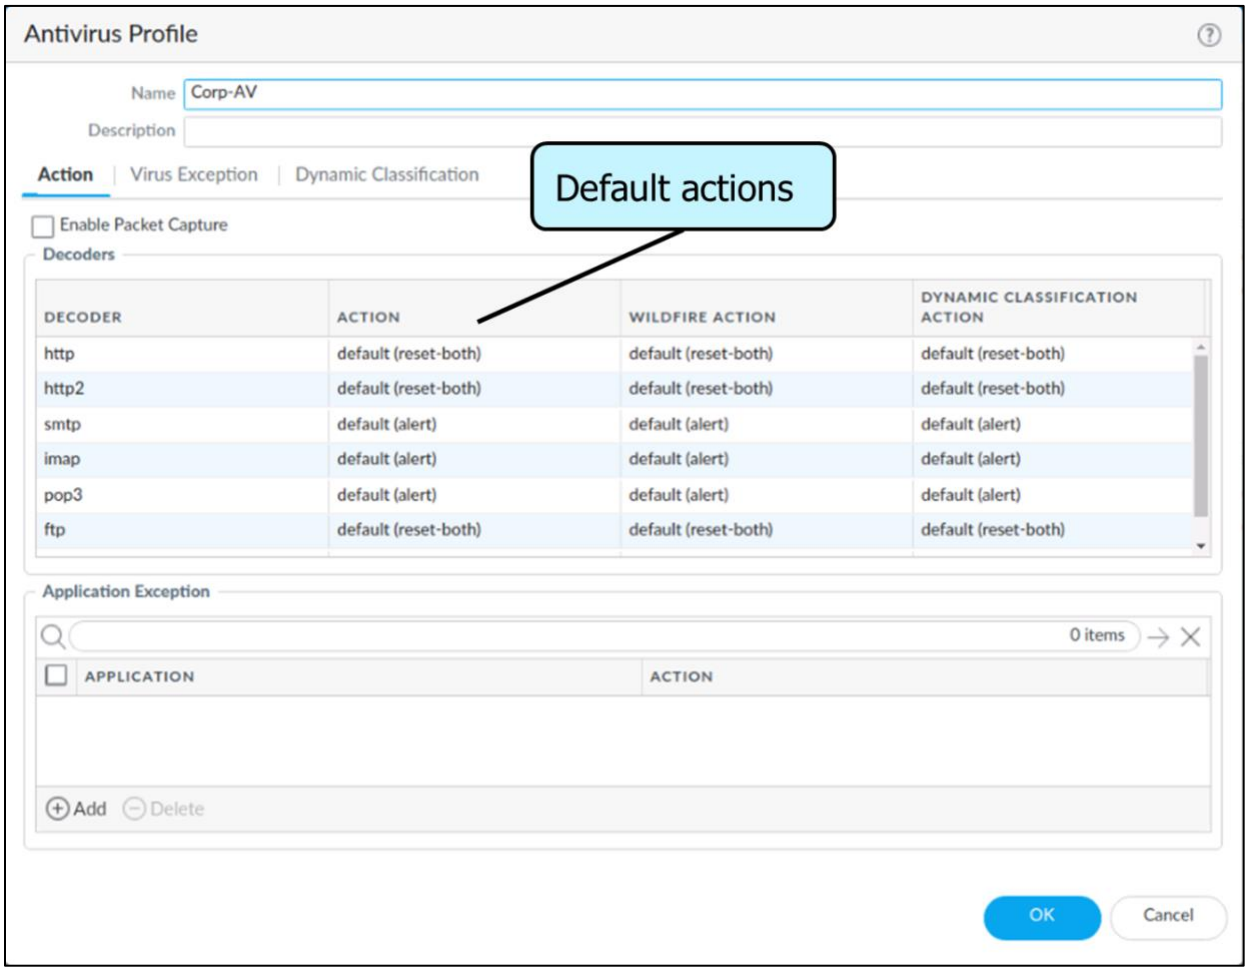

Anti-Spyware Security Profile Actions

You can create custom Anti-Spyware Profiles, or one of the two following predefined profiles can be chosen when anti-spyware is applied to a Security policy rule:

Profile	Description
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.

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

Action	Description
Default	For each threat signature and anti-spyware signature that is defined by Palo Alto Networks, a default action is specified internally. The default action typically is an "alert" or a "reset-both." The default action is displayed in parentheses, 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 the client and server ends. For UDP, drops the connection.
Block IP	Blocks traffic from either a source or a source-destination pair. It is configurable for a specified period of time.

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

Two Predefined Antivirus Security Profiles

	NAME	PACKET CAPTURE	Decoders				Dynamic Classification		THREAT EXCEPTIONS	DYNAMIC CLASSIFICATION EXCEPTIONS
			NAME	ACTION	WILDFIRE ACTION	DYNAMIC CLASSIFICATION ACTION	NAME	POLICY ACTION		
☐	default	☐	http	default (reset-both)	default (reset-both)	default (reset-both)	Windows Executables	enable	0	0
			http2	default (reset-both)	default (reset-both)	default (reset-both)	PowerShell Script 1	enable		
			smtp	default (alert)	default (alert)	default (alert)	PowerShell Script 2	enable		
			imap	default (alert)	default (alert)	default (alert)				
			pop3	default (alert)	default (alert)	default (alert)				
			ftp	default (reset-both)	default (reset-both)	default (reset-both)				
			smb	default (reset-both)	default (reset-both)	default (reset-both)				
☐	Corp-AV	☐	http	default (reset-both)	default (reset-both)	default (reset-both)	Windows Executables	enable	0	0
			http2	default (reset-both)	default (reset-both)	default (reset-both)	PowerShell Script 1	enable		
			smtp	default (alert)	default (alert)	default (alert)	PowerShell Script 2	enable		
			imap	default (alert)	default (alert)	default (alert)				
			pop3	default (alert)	default (alert)	default (alert)				
			ftp	default (reset-both)	default (reset-both)	default (reset-both)				
			smb	default (reset-both)	default (reset-both)	default (reset-both)				

Two default profiles

Vulnerability Protection Security Profile Actions

The *default* Vulnerability Protection Security Profile protects clients and servers from all known critical-, high-, and medium-severity threats. You also can create exceptions that enable you to change the response to a specific signature.

Two Predefined Vulnerability Protection Security Profiles

☐	NAME	COUNT	RULE NAME	THREAT NAME	HOST TYPE	SEVERITY	ACTION	PACKET CAPTURE
☐	strict	Rules: 10	simple-client-critical	any	client	critical	reset-both	disable
			simple-client-high	any	client	high	reset-both	disable
			simple-client-medium	any	client	medium	reset-both	disable
			simple-client-informational	any	client	informational	default	disable
			simple-client-low	any	client	low	default	disable
			simple-server-critical	any	server	critical	reset-both	disable
			simple-server-high	any	server	high	reset-both	disable
			more...					
☐	default	Rules: 6	simple-client-critical	any	client	critical	default	disable
			simple-client-high	any	client	high	default	disable
			simple-client-medium	any	client	medium	default	disable
			simple-server-critical	any	server	critical	default	disable
			simple-server-high	any	server	high	default	disable
			simple-server-medium	any	server	medium	default	disable

Two default profiles

URL Filtering Security Profile Actions

Action	Description
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 of 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 typically is used for categories that are considered benign and is used to improve the user experience by giving the user the option to continue if they consider the site to be 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 webpage 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 administrator or help-desk 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. The Override webpage doesn't display properly on client systems configured to use a proxy server.
none	The "none" action applies only to custom URL categories. Select **none** to ensure that, if multiple URL Filtering Profiles exist, the custom category will not have any impact on other profiles. For example, if you have two URL Filtering 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, to delete a custom URL category, the category must be set to **none** in any profile where it is used.

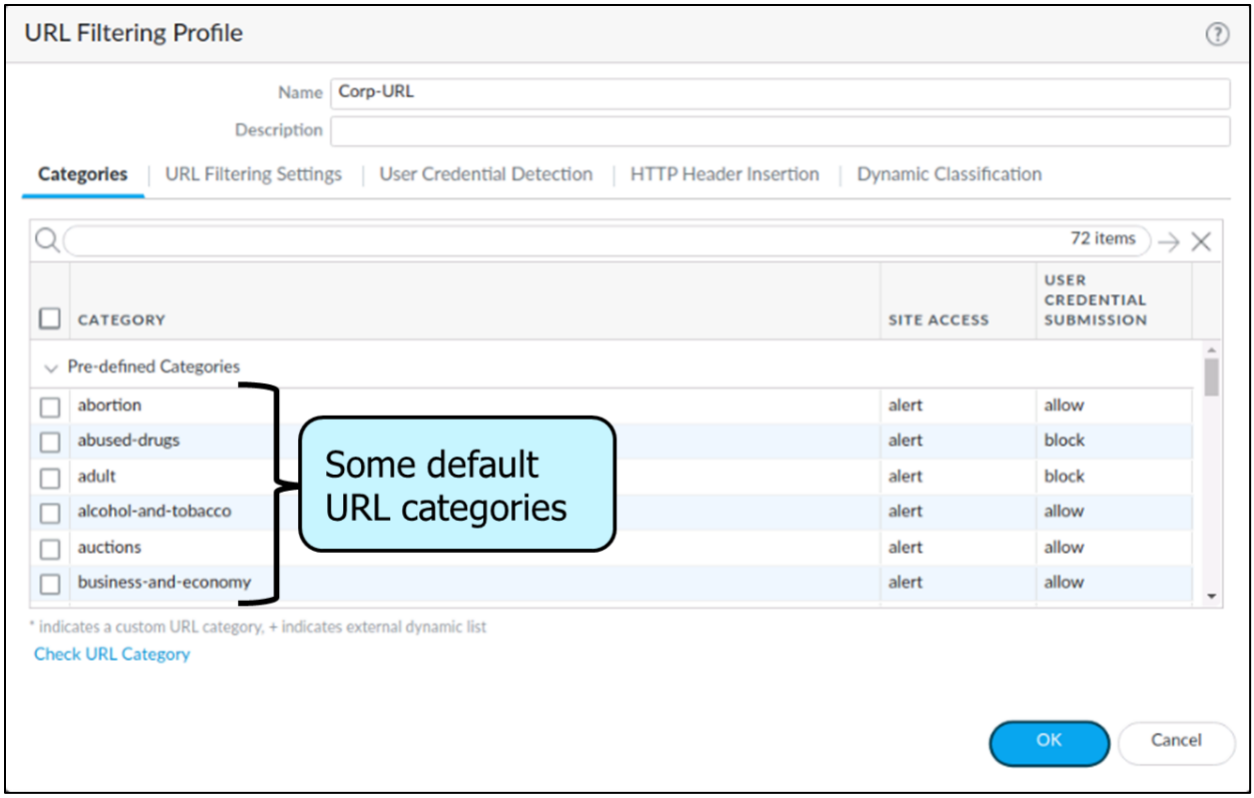

File Blocking Security Profile Actions

Field	Description
Name	Enter a rule name (up to 31 characters in length).
Applications	Select the applications the rule applies to or select **Any**.
File Types	Click in the field and then click **Add** to display a list of supported file types. Click a file type to add it to the profile and continue to add file types as needed. If you select **Any**, the defined action is taken on all supported file types.
Direction	Select the direction of the file transfer (**upload**, **download**, or **both**).
Action	Select the action taken when the selected file types are detected: • **alert:** An entry is added to the Threat log. • **block:** The file is blocked. • **continue:** A message to the user indicates that a download has been requested and asks the user to confirm whether to continue. The purpose is to warn the user of a possible unknown download (also known as a drive-by-download) and to give the user the option of continuing or stopping the download.

When you create a File Blocking Profile with the action “continue,” you can choose only the application web-browsing. If you choose any other application, traffic that matches the Security policy will not flow through the firewall because the users will not be prompted with a continue page.

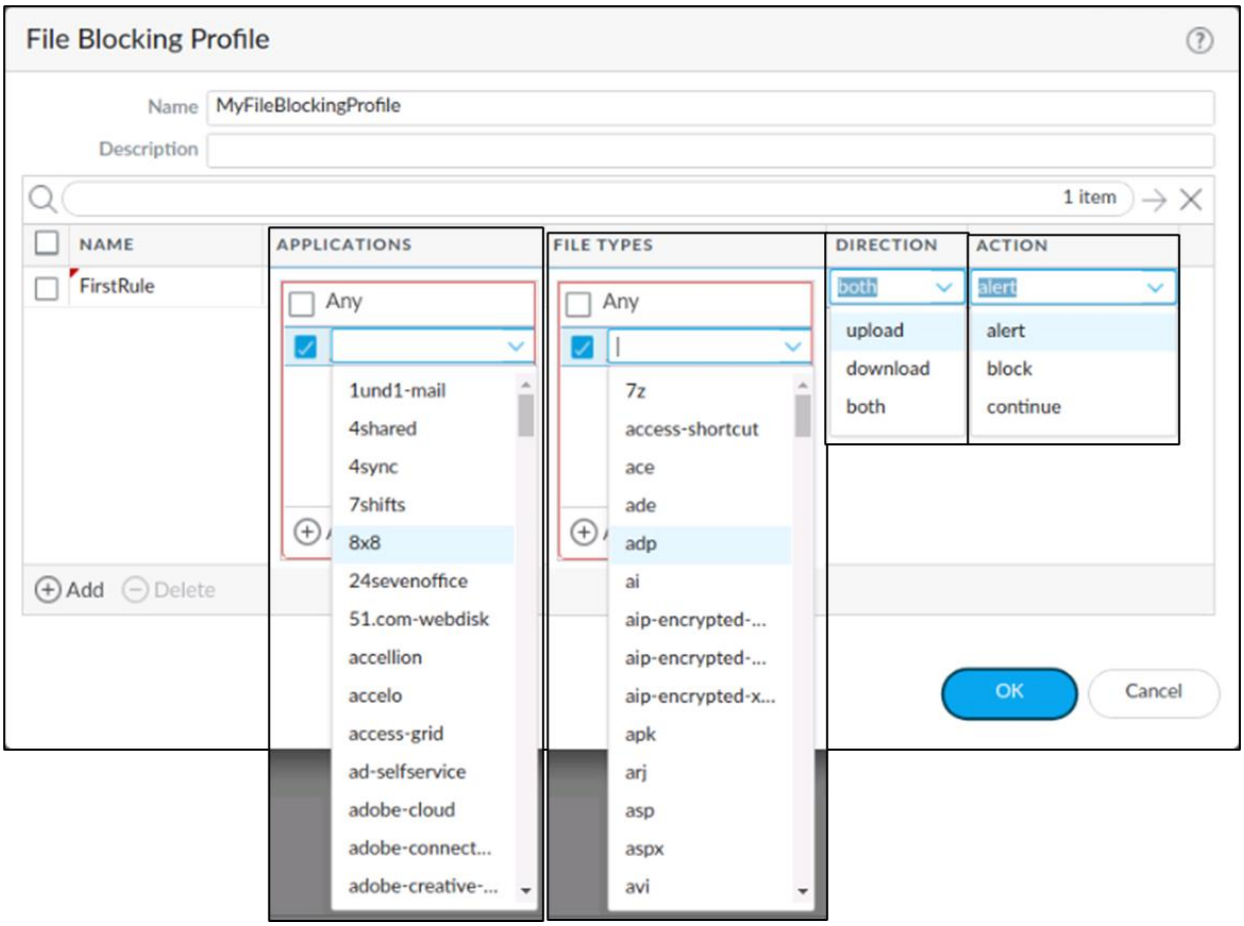

Sample Question

Q1. Which two actions are available for Antivirus Security Profiles? (Choose two.)

a) continue

b) allow

c) block IP

d) alert

Antivirus Security Profile Customization

Customized profiles can be used to minimize antivirus inspection for traffic between more trusted security zones. They also can be used to maximize the inspection of traffic received from less trusted zones such as the internet and of traffic sent to highly sensitive destinations such as server farms. The Palo Alto Networks WildFire product also provides signatures for persistent threats that are more evasive and have not yet been discovered by other antivirus solutions. As threats are discovered by WildFire, signatures are quickly created and then integrated into the standard antivirus signatures that can be downloaded daily by Threat Prevention subscribers (sub-hourly for WildFire subscribers).

Antivirus Security Profile Customization for Decoder Actions

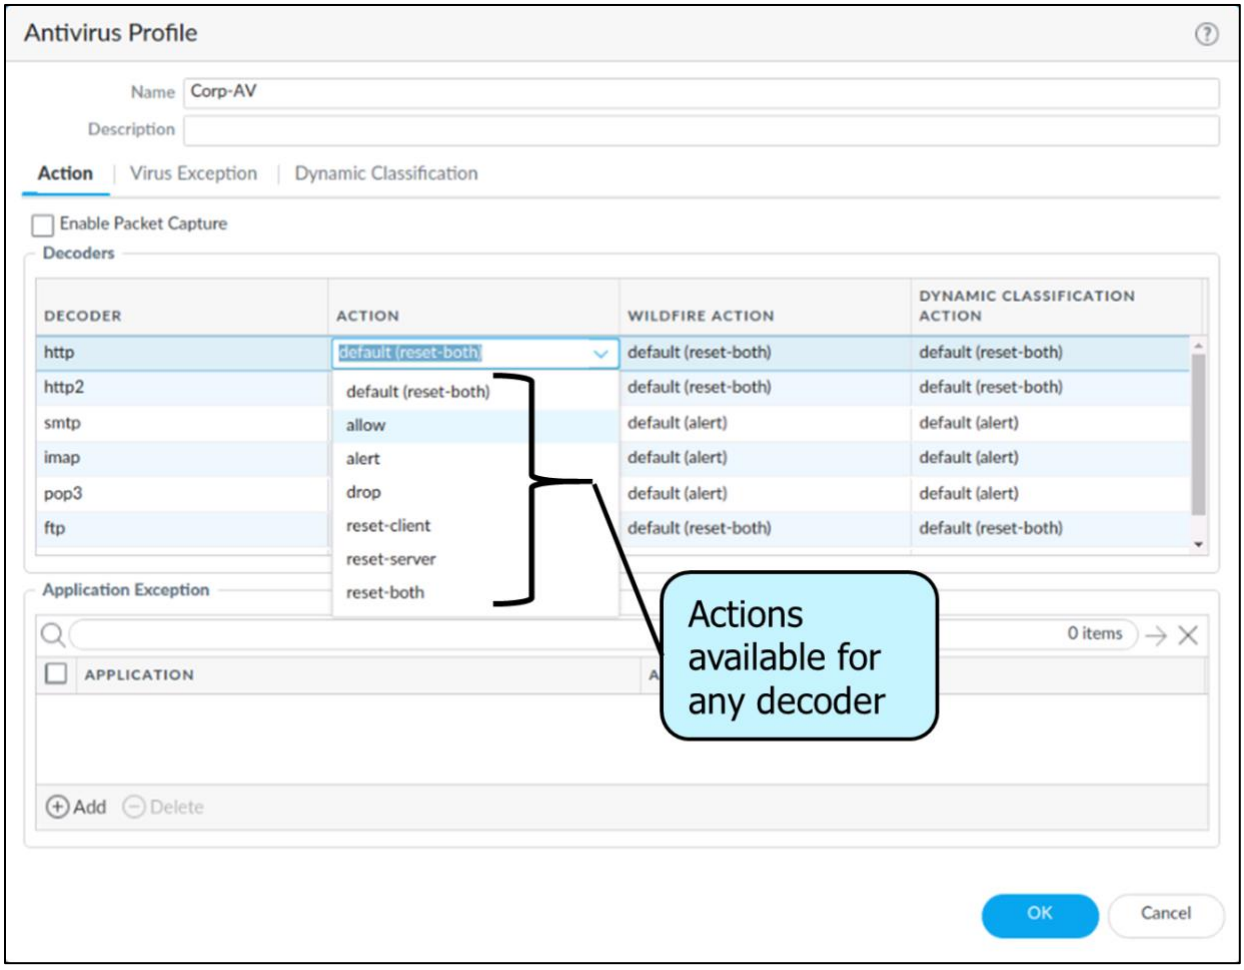

Antivirus Security Profile Customization to Exclude Specific Apps from AV Inspection

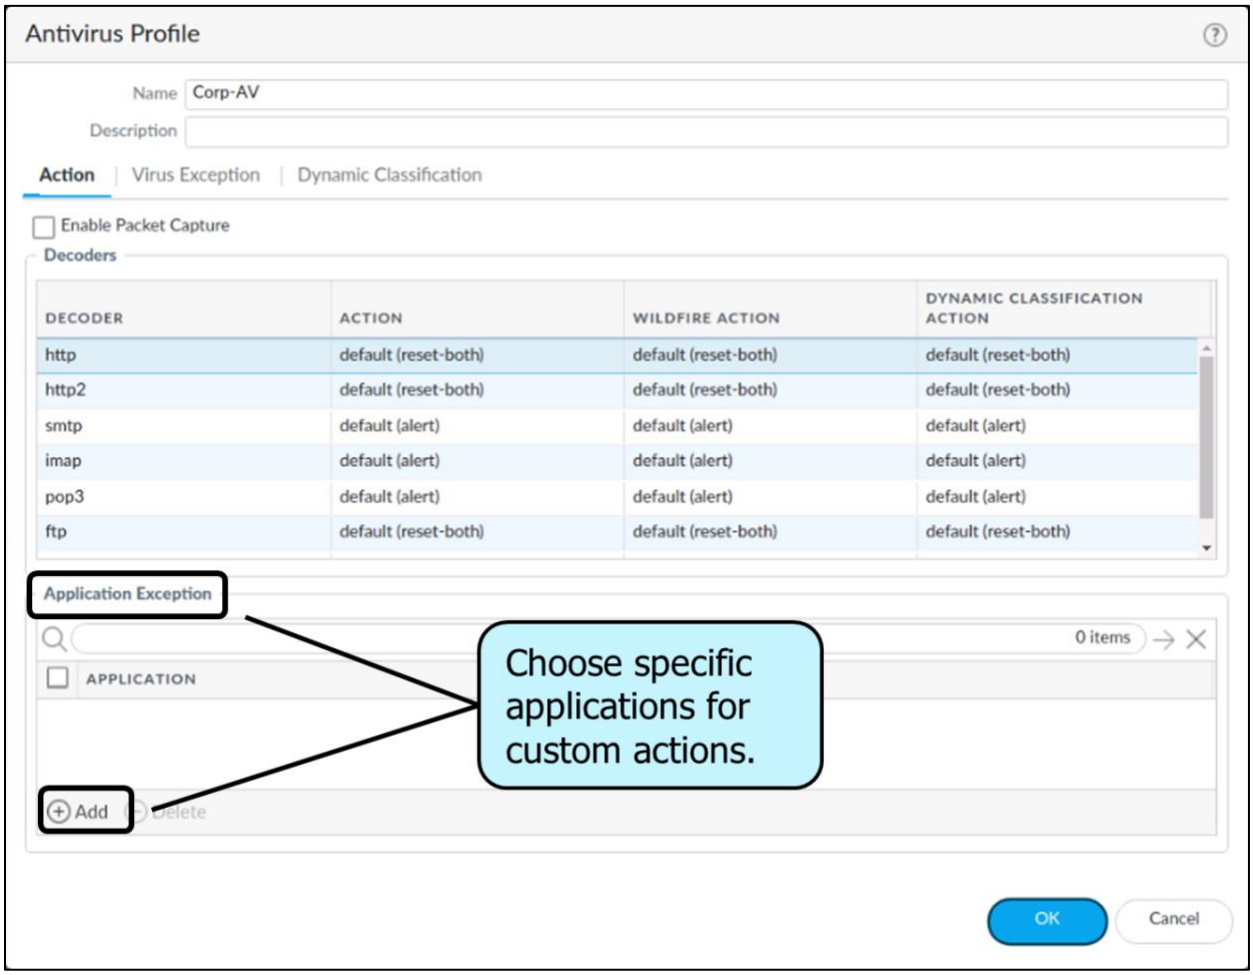

Antivirus Security Profile Customization to Exclude Threats from AV Inspection

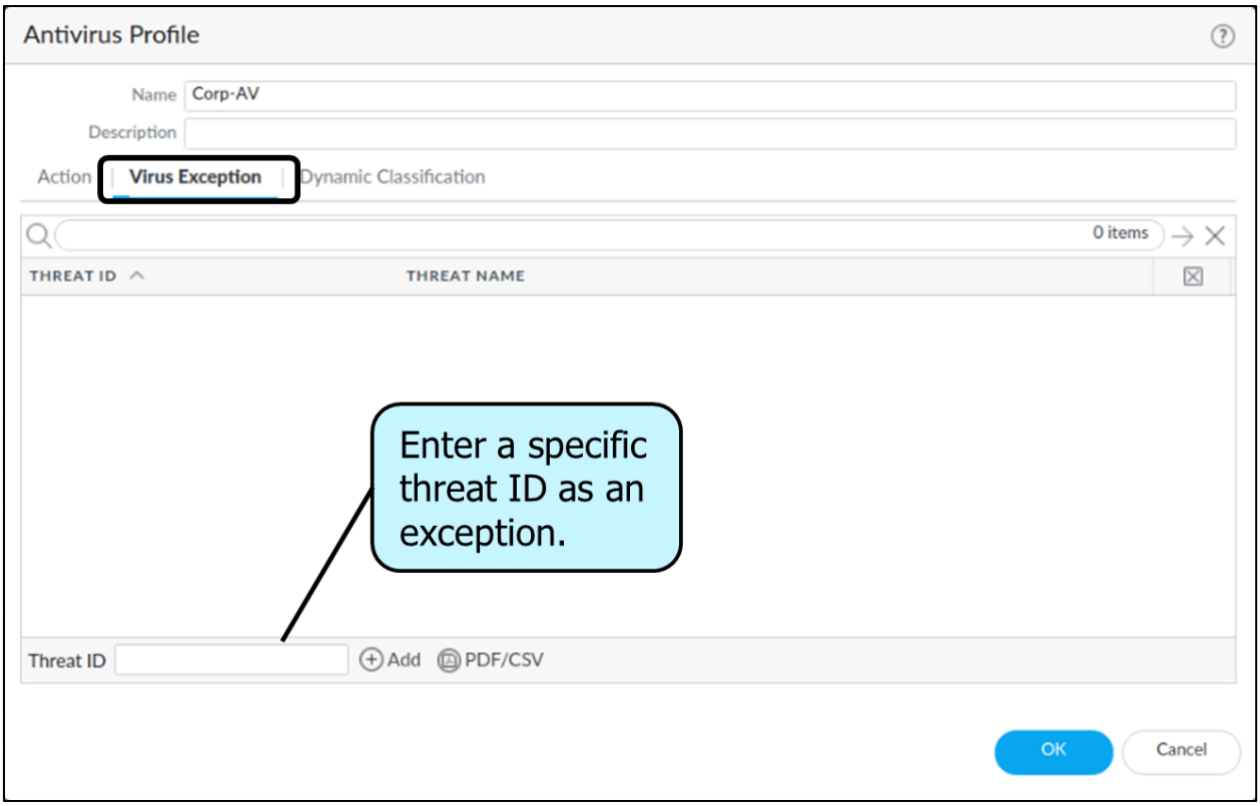

Anti-Spyware Security Profile Customization

Custom Anti-Spyware Profiles can be created. For example, you can reduce the stringency for anti-spyware inspection for traffic between more trusted security zones and maximize the inspection of traffic received from the internet or traffic sent to protected assets such as server farms.

Anti-Spyware Security Profile Customization of Categories

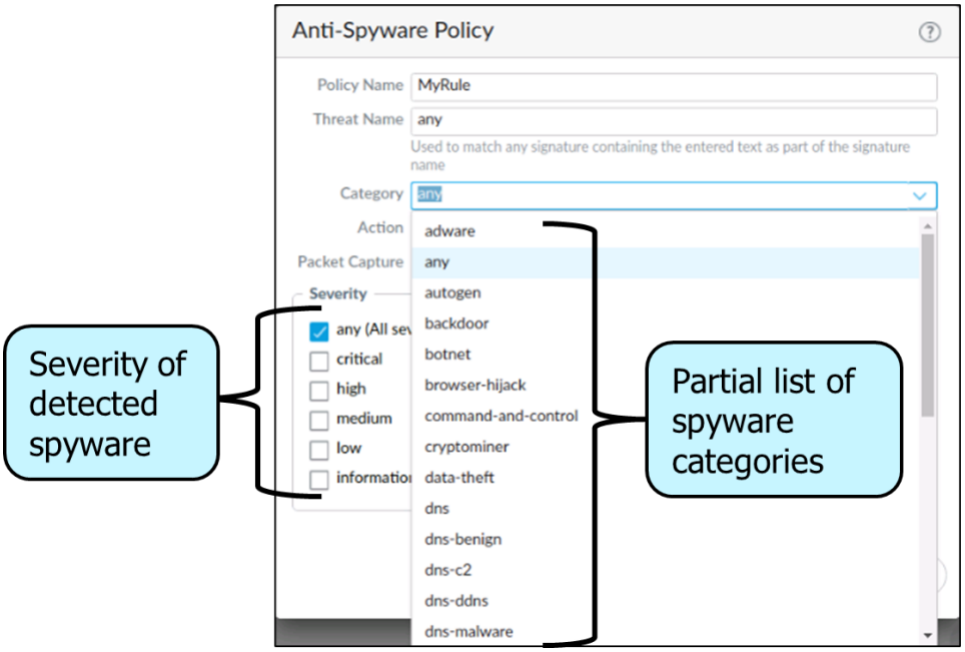

Anti-Spyware Security Profile Customization of Actions

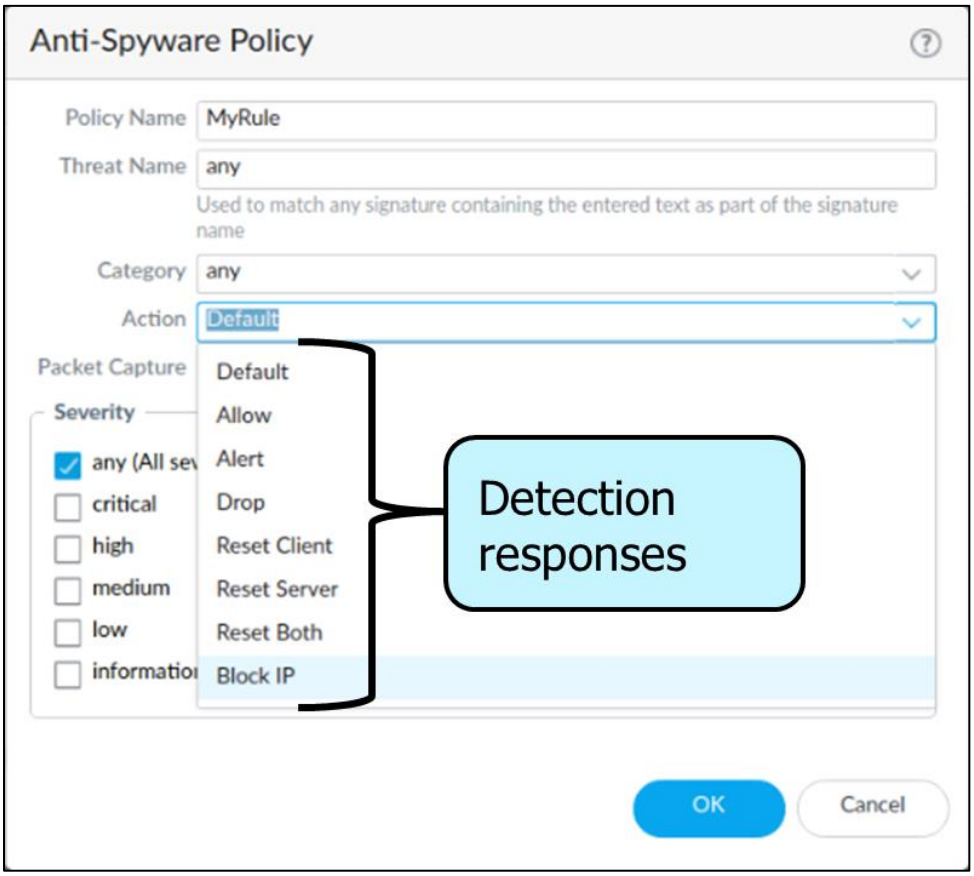

Vulnerability Protection Security Profile Customization

The *default* Vulnerability Protection Security Profile protects clients and servers from all known critical-, high-, and medium-severity threats. Customized profiles can be used to minimize vulnerability checking for traffic between more trusted security zones and to maximize protection for traffic received from less trusted zones such as the internet, along with traffic sent to highly sensitive destinations such as server farms.

The **Exceptions** setting found under the **Exceptions** tab enables you to change the response for a specific signature based on its Threat ID number or name. For example, you can block all packets that match specific signatures, except for the one(s) that you set up as exception(s), which could be set up as an action to generate only alerts.

Vulnerability Protection Profile Customization of CVEs and Vendor IDs

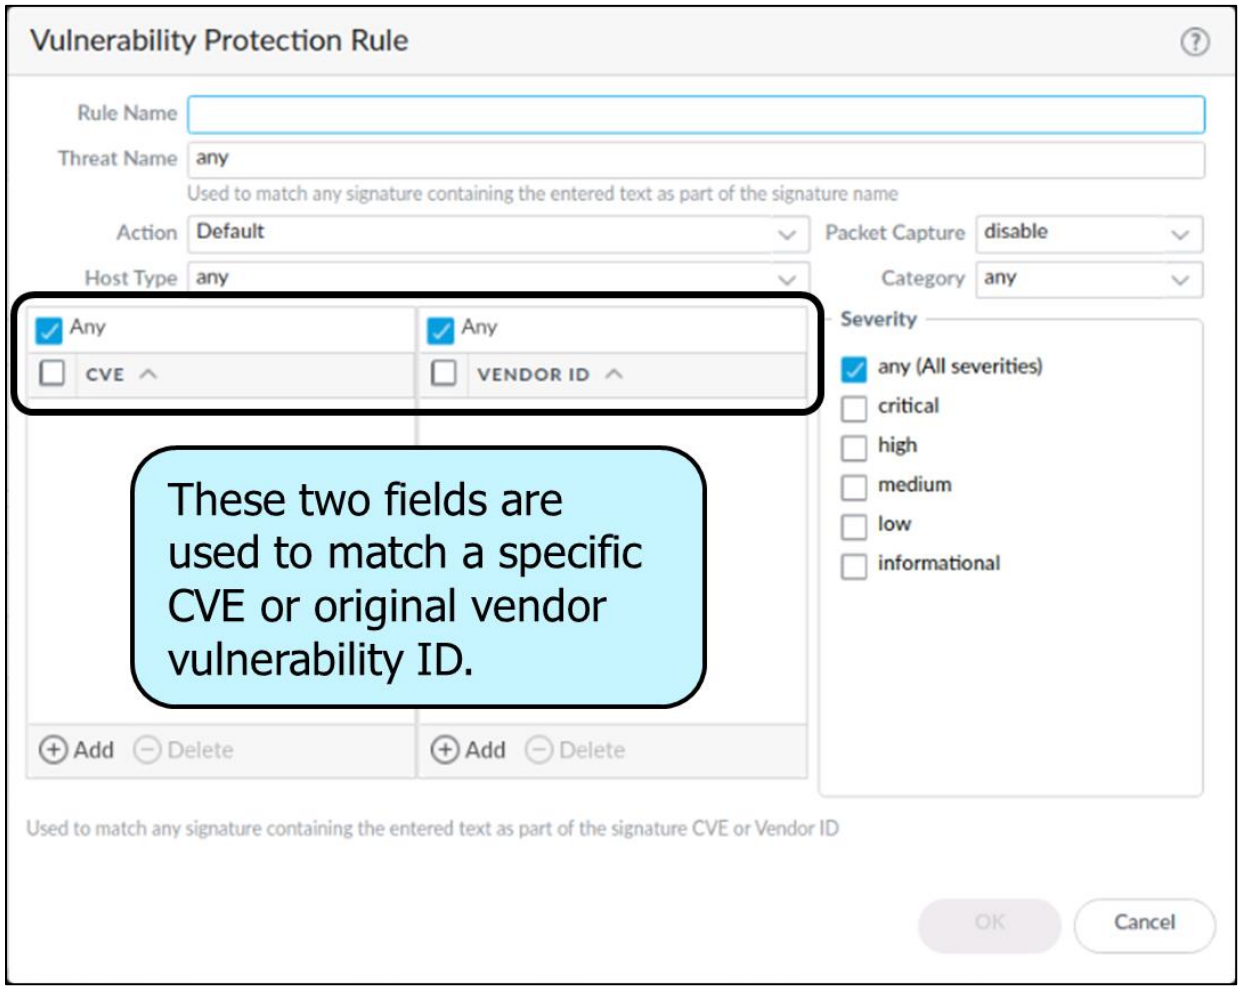

Vulnerability Protection Profile Customization of Actions and Severity

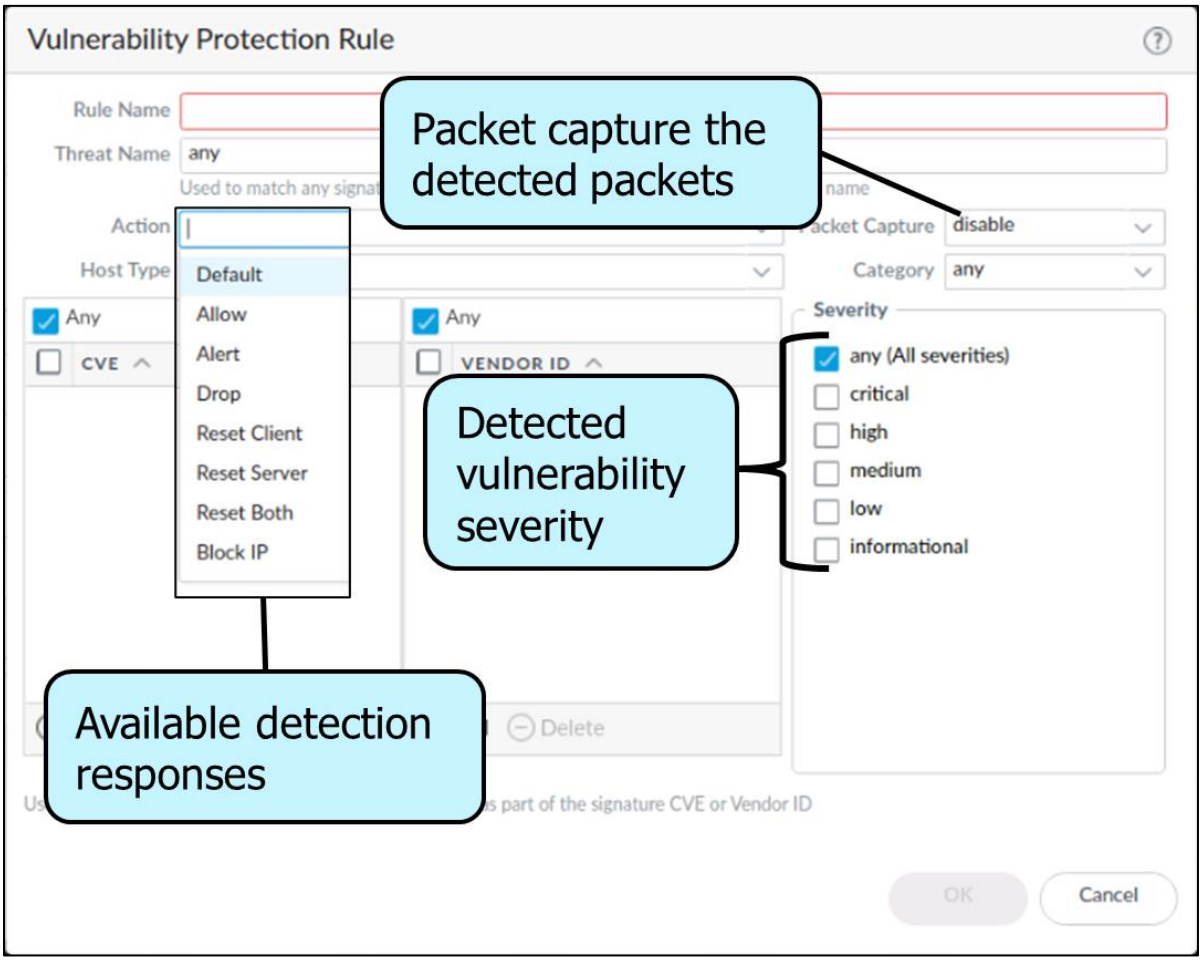

Vulnerability Protection Profile Customization of Categories

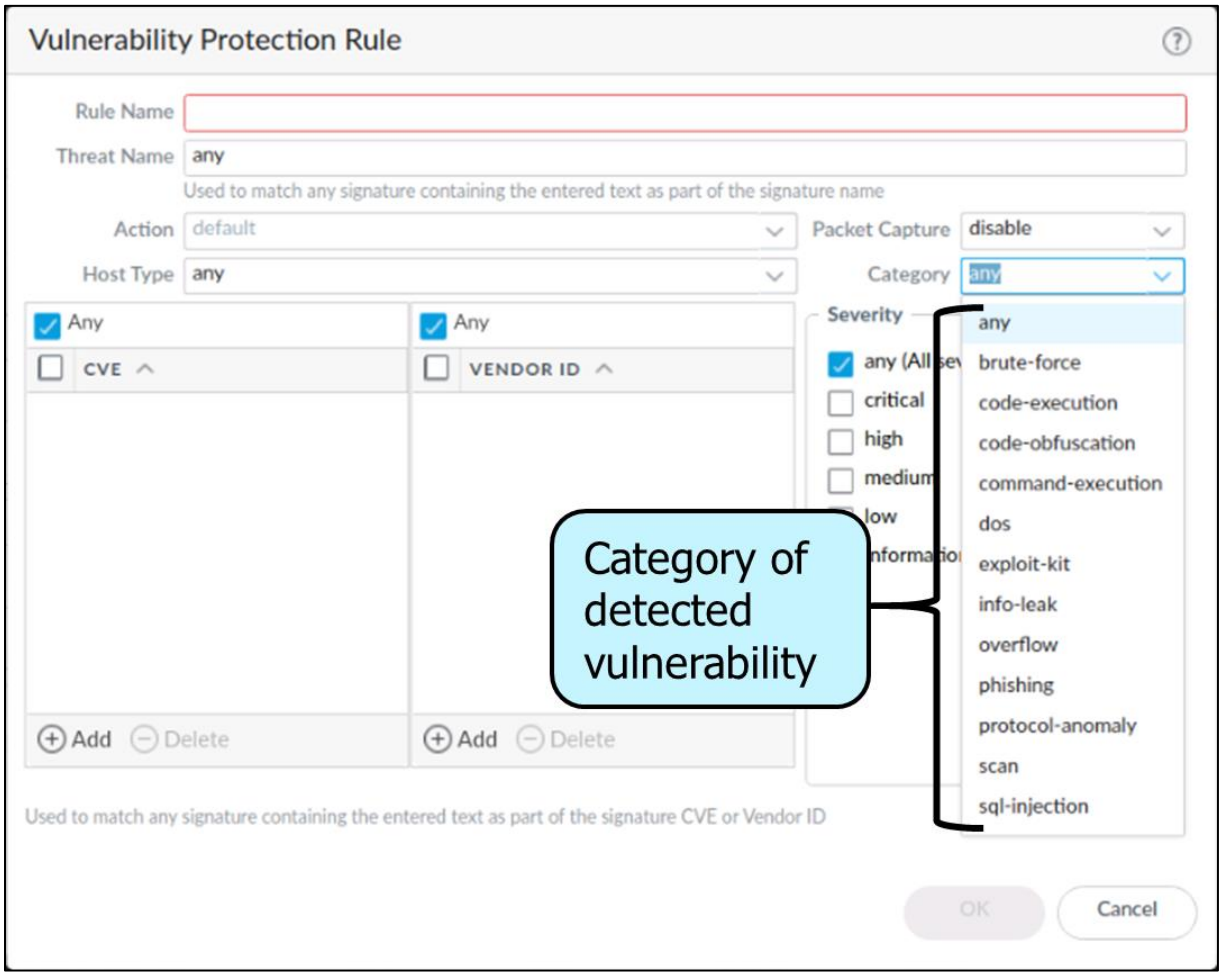

URL Filtering Security Profile Customization

URL Filtering Security Profiles should be customized to meet the unique needs of your organization.

URL Filtering Profile Customization of Categories and Site Access

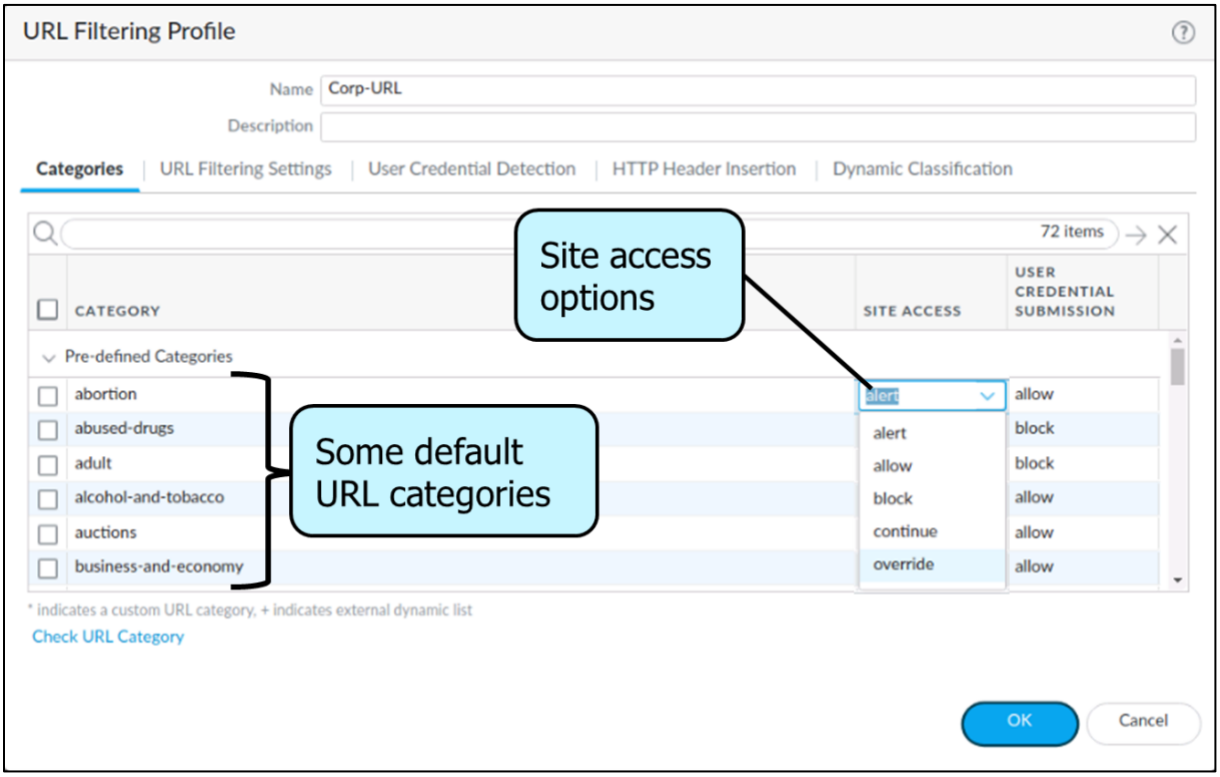

Safe Search

Many search engines have a safe search setting that filters out pornographic images and videos in search query return traffic. When **Safe Search Enforcement** is enabled, the firewall blocks search results if the end user is not using the strictest safe search settings in the search query. The firewall can enforce safe search for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. This is a best-effort setting and is not guaranteed by the search providers to work with every website.

HTTP Header Logging

The **HTTP Header Logging** feature provides visibility into the attributes included in the HTTP request sent to a server. When HTTP Header Logging is enabled, one or more of the following attributes are recorded in the URL Filtering log:

- **User Agent:** The web browser that the user used to access the URL. This information is sent in the HTTP request to the server. For example, the User Agent can be Internet Explorer or Firefox.
- **Referer:** The URL of the webpage that linked the user to another webpage. It is the source that redirected (referred) the user to the webpage that is being requested.
- **X-Forward-For**: The header field option that preserves the IP address of the user who requested the webpage. It enables you to identify the IP address of the user, which is particularly useful if you have a proxy server on you network or you have implemented source NAT that is masking the user's IP address such that all requests seem to originate from the proxy server's IP address or a common IP address.

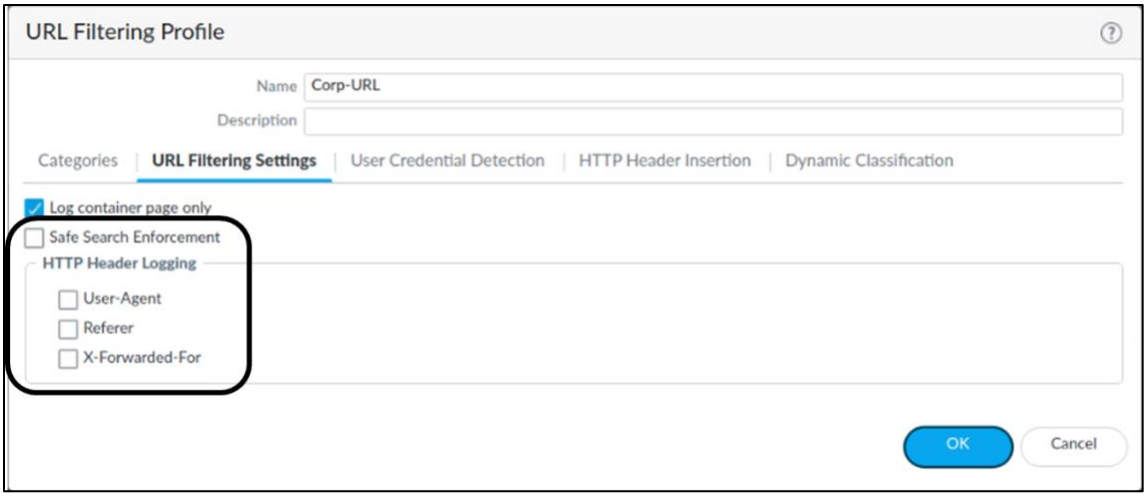

File Blocking Security Profile Customization

File blocking should be customized to meet the unique needs of your organization.

File Blocking Applications Customization

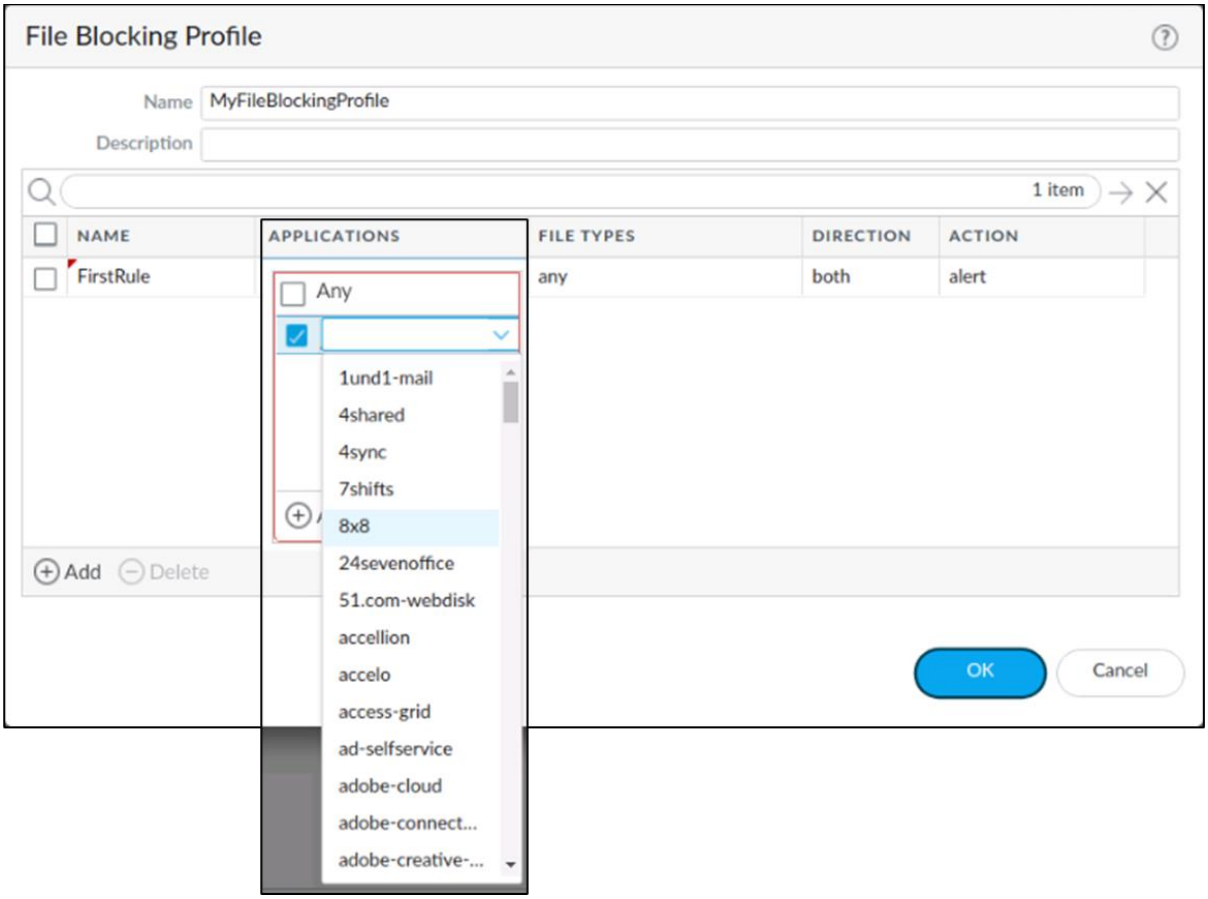

File Blocking File Types Customization

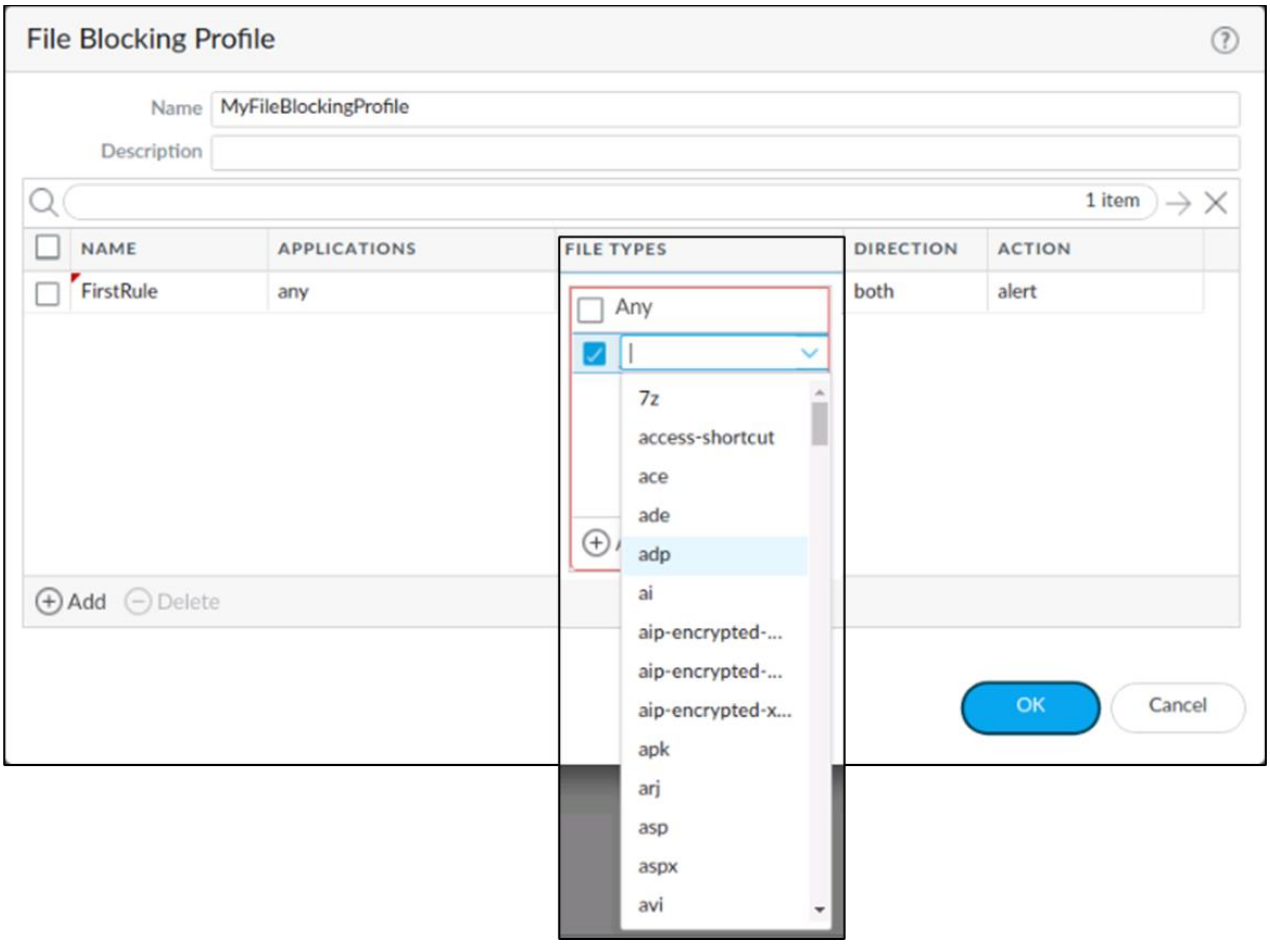

File Blocking Direction Customization

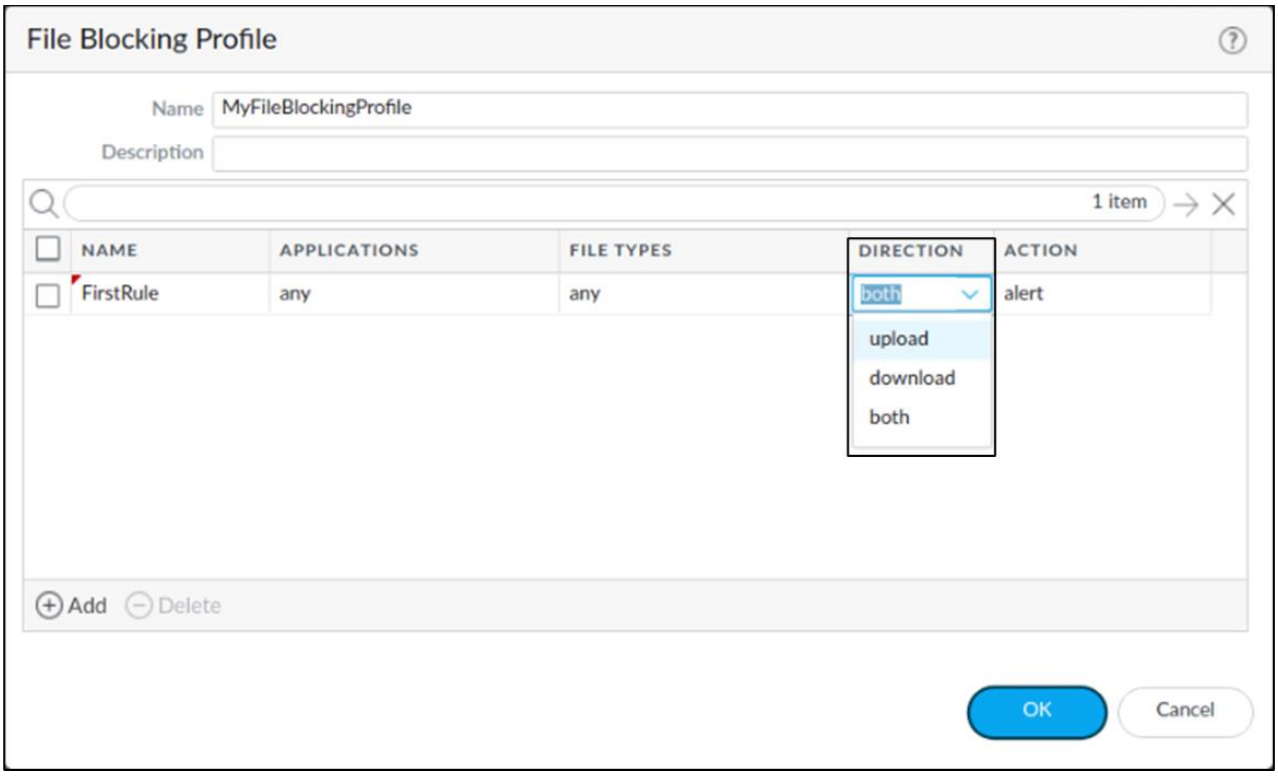

File Blocking Action Customization

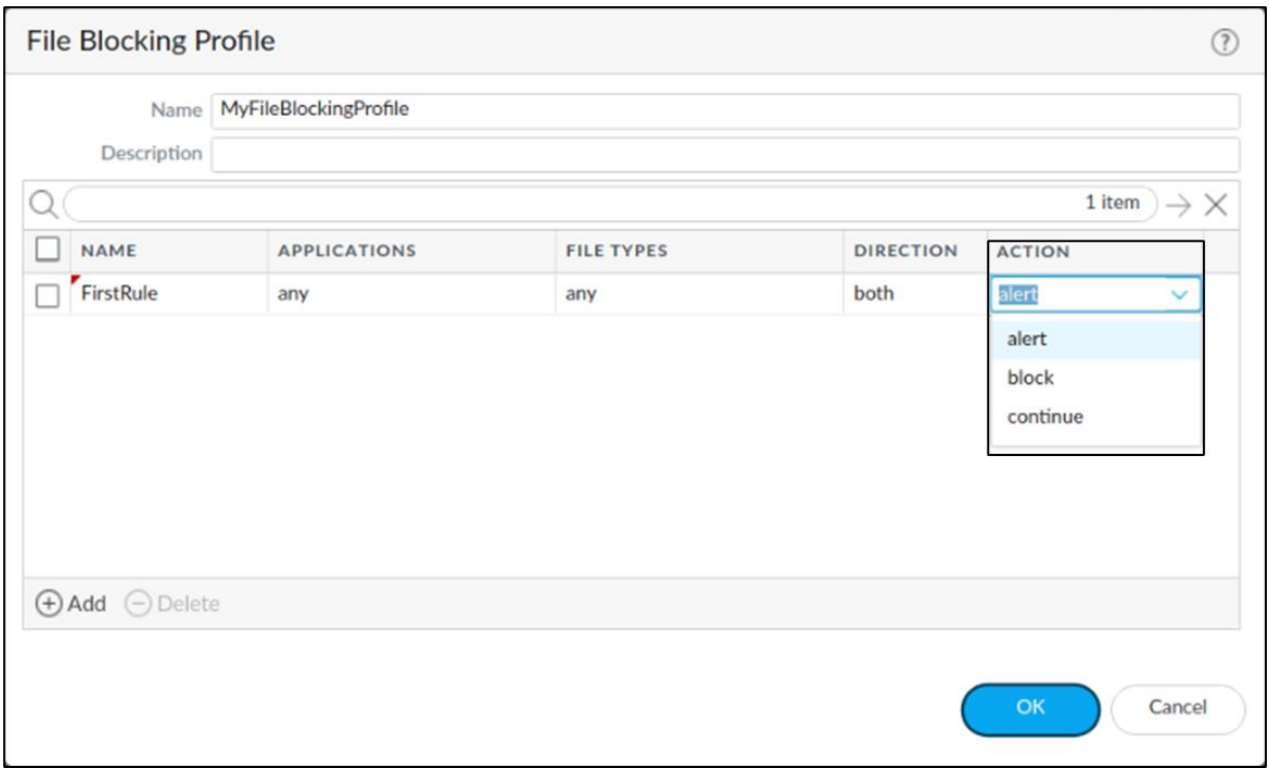

Sample Question

Q2.Which two HTTP Header Logging options are within a URL Filtering Profile? (Choose two.)

a) User-Agent

b) Safe Search

c) URL redirection

d) X-Forwarded-For

5.3 Use the cloud DNS Security to control traffic based on domains

DNS Security subscription enables users to access real-time protections using advanced predictive analytics. When techniques such as DGA/DNS tunneling detection and machine learning are used, threats hidden within DNS traffic can be proactively identified and shared through an infinitely scalable cloud service. Because the DNS signatures and protections are stored in a cloud-based architecture, you can access the full database of ever-expanding signatures that have been generated using a multitude of data sources. This list of signatures allows you to defend against an array of threats using DNS in real-time against newly generated malicious domains. To combat future threats, updates to the analysis, detection, and prevention capabilities of the DNS Security service will be available through content releases. To access the DNS Security service, you must have a Threat Prevention license and DNS Security license.

Enable DNS Security

To enable DNS security, domain queries using DNS security that are found to be threats are remediated with an Anti-Spyware Security Profile. Edit an existing or open a new Anti-Spyware Profile using **Objects > Security Profiles > Anti-Spyware**.

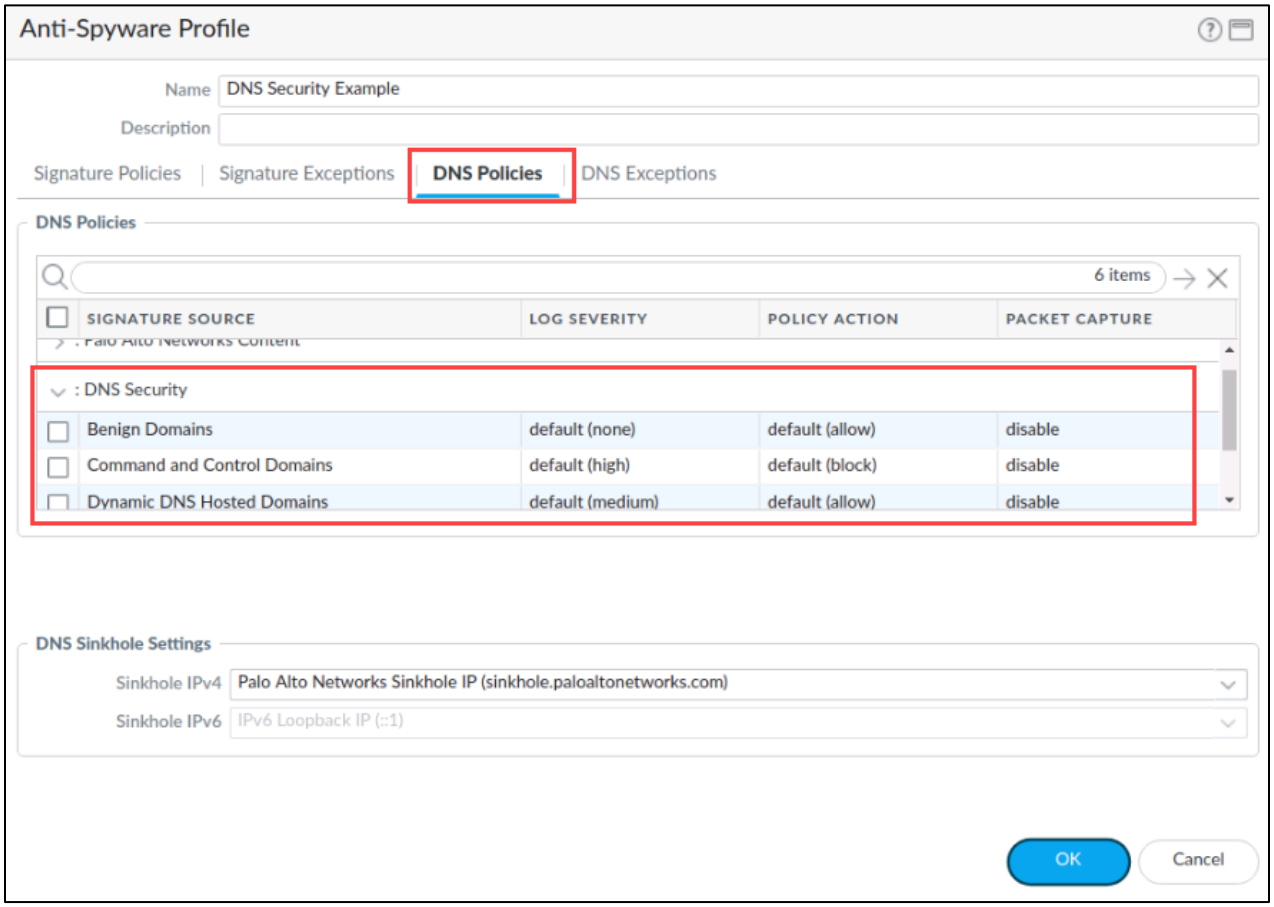

Click the **DNS Policies** tab and expand the **DNS Security** group item in the list of signature sources. Each list item is a cloud-based collection of DNS identifying information of the threat type indicated in the list item name. The **Policy Action** column shows the selected remediation when a threat is found in a list. An explanation of these actions, including Sinkholing, can be found in item 4.2 under Anti-Spyware Security Profile Actions. Anti-Spyware Profiles configured for DNS security protections are added to Security Profiles allowing traffic to be inspected.

Sample Question

Q1. Which two actions are required to implement DNS Security inspections of traffic? (Choose two.)

a) add an Anti-Spyware Security Profile with DNS remediations to a Security policy

b) enabled the Advanced DNS Security check box in General Settings

c) configure an Anti-Spyware Security Profile with DNS remediations

d) enter the address for the Secure DNS Service in the firewalls DNS settings

5.4 Use the PAN-DB database to control traffic based on websites

Most attacks and exposure to malicious content occur during normal web browsing activities, which means that all users must have safe, secure web access. PAN-DB is a global URL and IP database, designed to fulfill an enterprise's web security needs. URL filtering with PAN-DB automatically prevents attacks that leverage the web as an attack vector, including phishing links in emails, phishing sites, HTTP-based command and control, malicious sites, and pages that carry exploit kits.

Granular policy enables the prevention of downloads, the automation of warning messages, or the restriction of access altogether. PAN-DB provides real-time protections from emerging attacks. PAN-DB receives updates from WildFire every five minutes to block malicious sites. There also are other advanced identification techniques.

PAN-DB is tightly integrated into PAN-OS software, thus providing advanced persistent threat (APT) protection with high performance beyond traditional URL filtering. Traditional URL filtering is intended to control unwanted web surfing such as non-business or illegal sites, but it usually doesn't cover up-to-the-minute malicious websites such as a newly discovered malware site, an exploit site, or command-and-control (C2) sites.

How PAN-DB Maximizes URL Lookup Performance

The following sections describe the components shown in the following figure.

PAN-DB Classification and Cache System

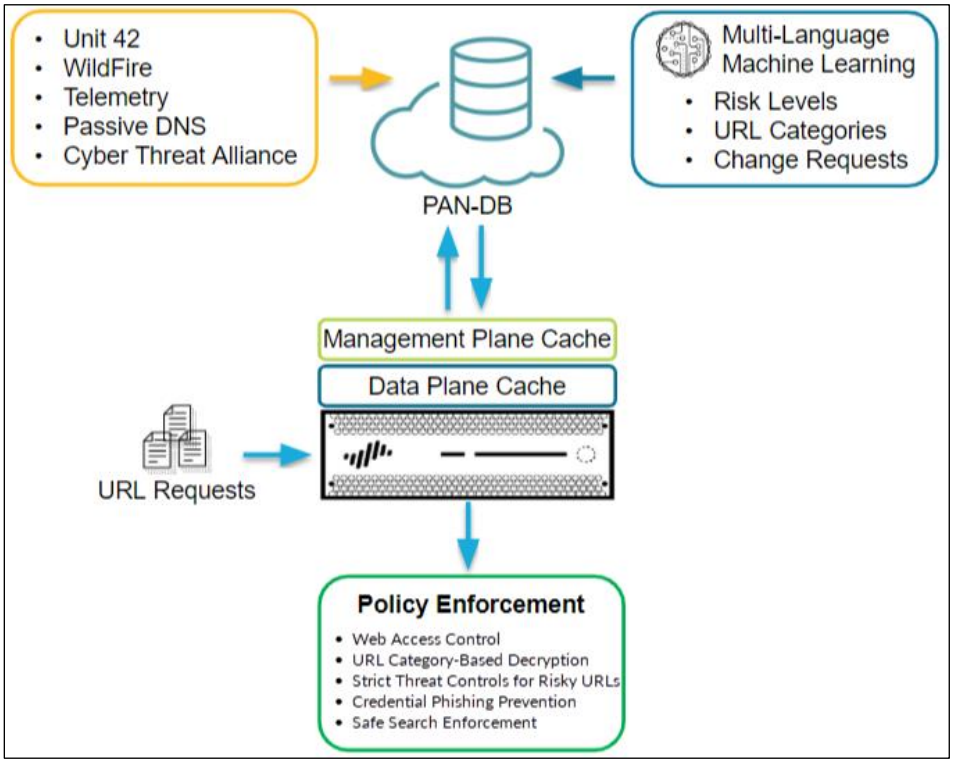

PAN-DB Core

The PAN-DB core, located in the Palo Alto Networks threat intelligence cloud, has a full URL and IP database to cover web security needs.

Management Plane Cache

The PAN-DB database is placed into the management plane cache, which provides quick URL lookups. The management plane cache will pull more URLs and categories from the PAN-DB core as users access sites that are not currently in the management plane cache. Any URL requested by a user that is "unknown" to Palo Alto Networks will be examined, categorized, and implemented as appropriate.

Data-Plane Cache

A data-plane cache contains the most frequently accessed sites, which enables quicker URL lookups. The Malicious URL database is delivered from WildFire.

Millions of URLs and IPs are classified in a variety of ways. The PAN-DB receives URLs and IP addresses from the "Multi-Language Classification Engine" and from "URL Change Requests from users." The PAN-DB also receives malicious URL and IP information from WildFire. Examples of malicious URLs and the IP database follow:

- Malware Download URL and IP address: Prevent from downloading malware
- C&C URL and IP address: Disable malware communications

The malicious URLs are generated as WildFire identifies unknown malware, zero-day exploits, and APTs by executing them in a virtual sandbox environment.

PAN-DB will block a malicious URL with low latency.

PAN-DB has a superior mechanism that increases the speed of URL lookups, which means that you will get URL category information without sacrificing throughput.

The malicious URLs are generated as WildFire identifies unknown malware, zero-day exploits, and APTs, and executes them in a virtual sandbox environment. The ongoing malicious URL updates to PAN-DB allows you to block malware downloads and disable malware C2 communications.

Use the malicious URL database to block access to a variety of malicious websites without compromising web access performance.

Sample Question

Q1. Which two types of attacks does the PAN-DB prevent? (Choose two.)

a) phishing sites

b) HTTP-based command and control

c) infected JavaScript

d) flood attacks

5.5 Identify how to control access to specific URLs using custom URL filtering categories

Custom URL Filtering Categories

Use the **Custom URL Category** page to create your custom list of URLs and use it in a URL Filtering Profile or as a match criterion in policy rules. In a custom URL category, you can add URL entries individually or import a text file that contains a list of URLs. URL entries added to custom categories are case-insensitive.

Custom URL category settings are as follows:

- **Name:** Enter a name to identify the custom URL category (up to 31 characters in length). This name displays in the category list when URL Filtering Profiles are defined and in the match criteria for URL categories in policy rules. The name is case-sensitive and must be unique. Use only letters, numbers, spaces, hyphens, and underscores.
- **Description:** Enter a description for the URL category (up to 255 characters in length).
- Sites:
    - Click **Add** to enter URLs, only one in each row. Each URL can be in the format "www.example.com" or can include wildcards ("*.example.com").
    - Click **Import** and browse to select the text file that contains the list of URLs. Enter only one URL per row. Each URL can be in the format "www.example.com" or can include wildcards ("*.example.com").
    - Click **Export** to export the custom URL entries included in the list. The URLs are exported as a text file.
    - Select an entry and click **Delete** to remove the URL from the list. Before you can delete a custom category that you have used in a URL Filtering Profile, you must set the action to **None**. Go to **Category** actions in **Objects > Security Profiles > URL Filtering**.

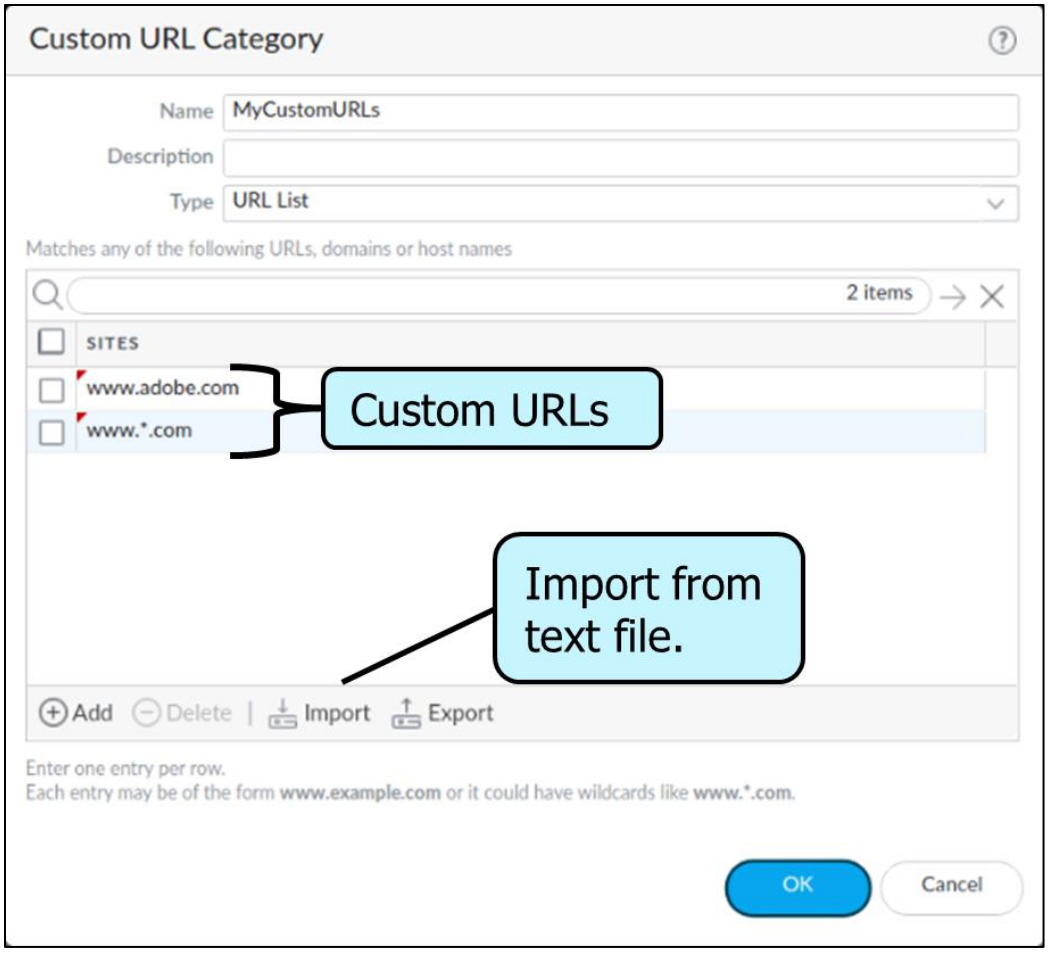
Custom URL Category
Name MyCustomURLs
Description
Type URL List
Matches any of the following URLs, domains or host names
2 items
SITES
www.adobe.com
www.*.com
Custom URLs
Import from text file.
Add
Delete
Import
Export
Enter one entry per row.
Each entry may be of the form www.example.com or it could have wildcards like www.*.com.
OK
Cancel

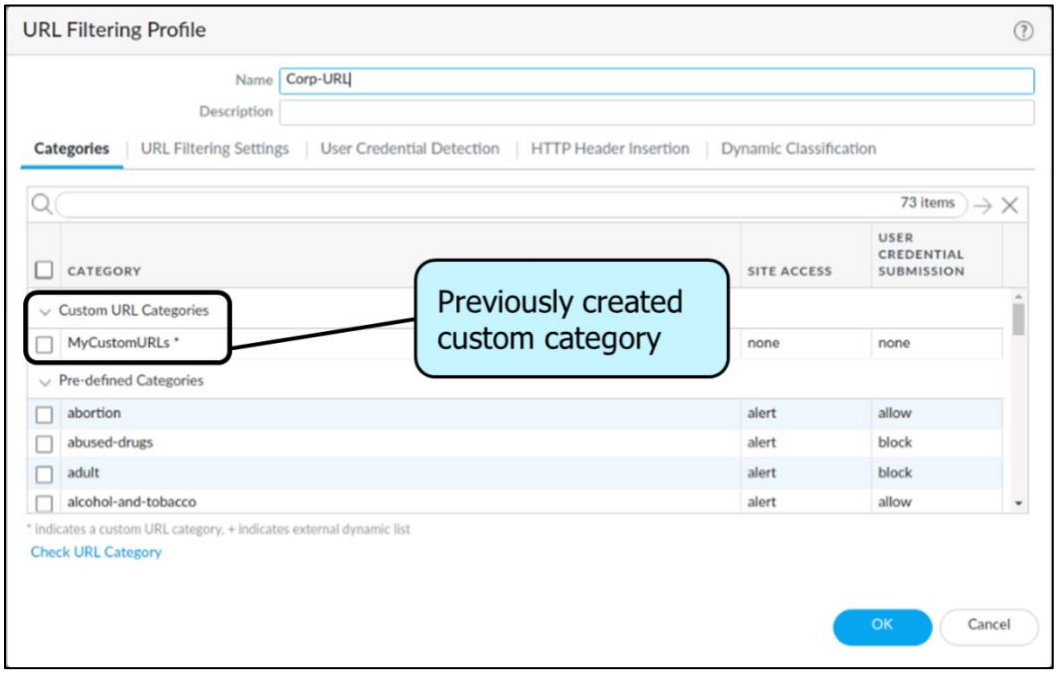
URL Filtering Profile
Name Corp-URL
Description
Categories
URL Filtering Settings
User Credential Detection
HTTP Header Insertion
Dynamic Classification
73 items
CATEGORY
SITE ACCESS
USER CREDENTIAL SUBMISSION
Custom URL Categories
MyCustomURLs *
none
none
Previously created custom category
Pre-defined Categories
abortion
alert
allow
abused-drugs
alert
block
adult
alert
block
alcohol-and-tobacco
alert
allow
* indicates a custom URL category, + indicates external dynamic list
Check URL Category
OK
Cancel

Sample Question

Q1. Which two valid URLs can be used in a custom URL category? (Choose two.)

a) ww.youtube.**

b) www.**.com

c) www.youtube.com

d) *youtube*

e) *.youtube.com

5.6 Differentiate between group mapping and IP to user mapping within policies and logs

Several options must be configured before User-ID can function. The LDAP Server Profile is the most important item to configure. The LDAP Server Profile is used to connect the firewall to an LDAP server and retrieve a list of usernames and groups.

The LDAP Server Profile will require different information, depending on what is used.

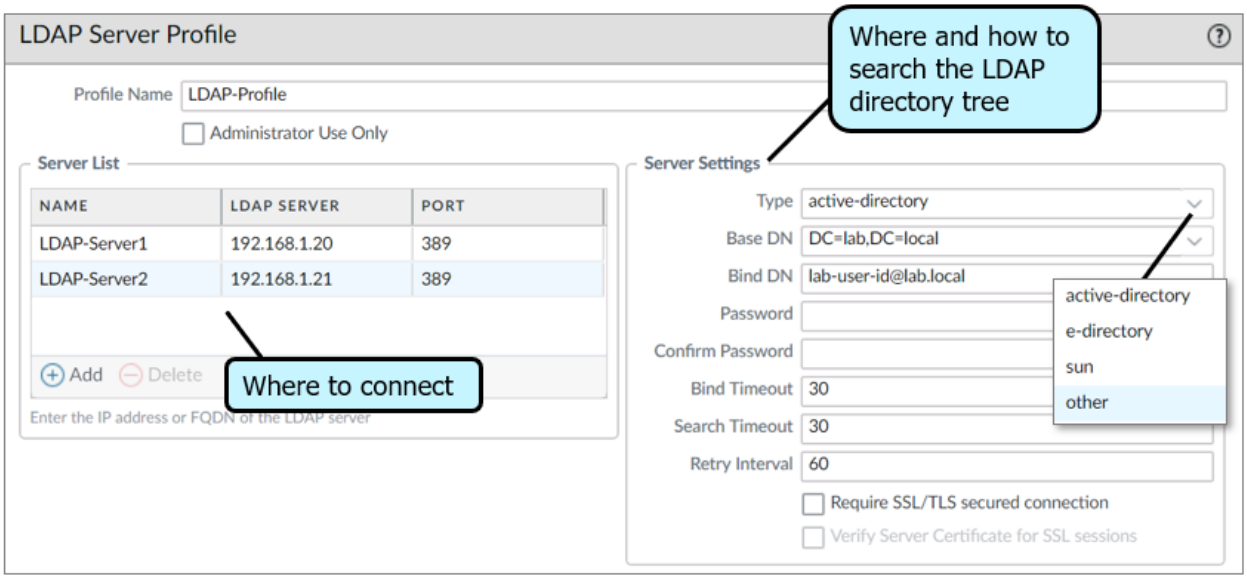

After the LDAP Server Profile is configured, the group mapping needs to be configured:

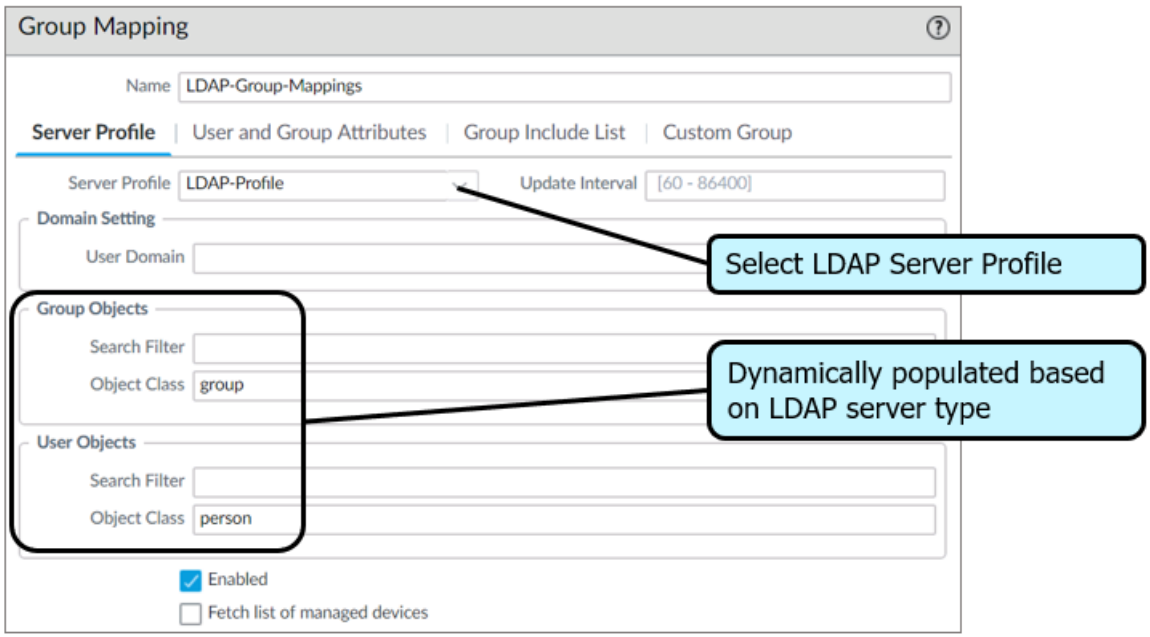

Administrators should select the LDAP Server Profile they configured earlier and complete the domain settings. The **Group Include List** tab will show the available groups in the domain. The administrator can choose which group to monitor and which ones to ignore:

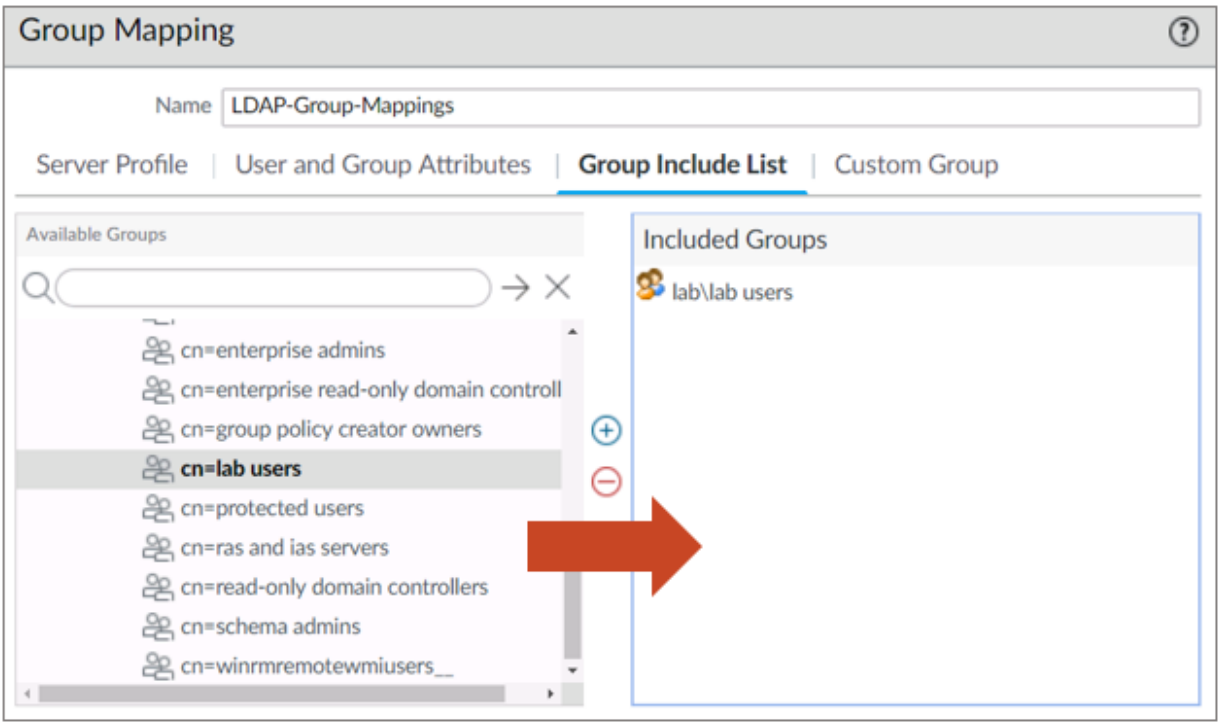

To learn more about the methods to map users and groups to collect User-ID information, see the following information:

- "Block Threats by Identifying Users " module in the EDU-210 training, Firewall Essentials: Configuration and Management
- User-ID in the *PAN-OS Administrator's Guide*

Answers to the Sample Questions

Correct answers are indicated in **bold**.

Domain 1 – Palo Alto Networks Strata Core Components

1.1 Understand the components of the Palo Alto Networks Strata Portfolio

Q1. What are four components of the Palo Alto Networks Cybersecurity Portfolio? (Choose four.)

a) Cortex DynamicDNS

b) WildFire

c) Cortex XDR

d) OpenConnect

e) Prisma Access

f) AutoFocus

Q2. Which cloud-delivered security service provides instant access to community-based threat data?

a) Prisma SaaS

b) AutoFocus

c) Unit 42

d) Cortex XDR

Q3. Which cloud-delivered security service provides security and connectivity for branches and mobile users?

a) Cortex XSOAR

b) Cortex XDR

c) AutoFocus

d) Prisma Access

Q4. Which Palo Alto Networks cybersecurity portfolio product provides access to applications from Palo Alto Networks, third parties, and customers?

a) WildFire

b) Cortex Data Lake

c) Network Security

d) Prisma Access

Q5. Which Palo Alto Networks firewall feature provides all the following abilities?

- Stops malware, exploits, and ransomware before they can compromise endpoints
- Provides protection while endpoints are online and offline, on network and off
- Coordinates enforcement with network and cloud security to prevent successful attacks
- Detects threats and automates containment to minimize impact
- Creates zero-day malware signatures with cloud-based threat analysis
- Integrates with Palo Alto Networks Cortex Data Lake

a) Cortex XDR

b) Prisma SaaS

c) WildFire

d) AutoFocus

1.2 Identify the components and operation of Single-Pass Parallel Processing architecture

Q1. Which three management features does the control plane provide? (Choose three.)

a) security processing

b) logging

c) reporting

d) firewall configuration

e) signature matching

f) network processing

Q2. Which three data processing features does the data plane provide? (Choose three.)

a) security processing

b) logging

c) reporting

d) firewall configuration

e) signature matching

f) network processing

Q3. What are three components of the Network Processing module? (Choose three.)

a) QoS

b) NAT

c) App-ID

d) flow control

e) url match

f) spyware

Q4. Which approach most accurately defines the Palo Alto Networks SP3 architecture?

a) prioritize first

b) sequential processing

c) scan it all, scan it once

d) Zero Trust segmentation platform

Q5. What is the result of using a stream-based architectural design?

a) superior performance

b) increased latency

c) detailed logging

d) increased functionality

Domain 2 – Device Management and Services

2.1 Identify and use firewall management interfaces

Q1. What are two firewall management methods? (Choose two.)

a) CLI

b) RDP

c) VPN

d) XML API

Q2. Which two devices are used to connect a computer to the firewall for management purposes? (Choose two.)

a) rollover cable

b) serial cable

c) RJ-45 Ethernet cable

d) USB cable

Q3. What is the default IP address on the MGT interfaces of a Palo Alto Networks firewall?

a) 192.168.1.1

b) 192.168.1.254

c) 10.0.0.1

d) 10.0.0.254

Q4. What are the two default services that are available on the MGT interface? (Choose two.)

a) HTTPS

b) SSH

c) HTTP

d) Telnet

Q5. True or false. Service route traffic has Security policy rules applied against it.

a) true

b) false

Q6. Service routes may be used to forward which two traffic types out a data port? (Choose two.)

a) External Dynamic Lists

b) MineMeld

c) Skype

d) Palo Alto Networks updates

2.3 Define firewall configurations

Q1. Which firewall plane does the running configuration reside on?

a) management

b) control

c) data

d) security

Q2. Which firewall plane does the candidate configuration reside on?

a) management

b) control

c) data

d) security

Q3. Candidate config and running config files are saved as which file type?

a) TXT

b) HTML

c) XML

d) RAR

Q4. Which command must be performed on the firewall to implement any changes?

a) commit

b) save

c) load

d) import

Q5. Which command backs up configuration files to a remote network device?

a) import

b) load

c) copy

d) export

Q6. The command **`load named configuration snapshot`** overwrites the current candidate configuration with which three items? (Choose three.)

a) custom-named candidate configuration snapshot (instead of the default snapshot)

b) custom-named running configuration that you imported

c) snapshot.xml

d) current running configuration (running-config.xml)

e) Palo Alto Networks updates

2.5 Identify the types of dynamic updates and their purpose

Q1. True or false. A Palo Alto Networks firewall automatically provides a backup of the configuration during a software upgrade.

a) true

b) false

Q2. If you have a Threat Prevention subscription but not a WildFire subscription, how long must you wait for the WildFire signatures to be added into the antivirus update?

a) 1 to 2 hours

b) 2 to 4 hours

c) 10 to 12 hours

d) 24 to 48 hours

Q3. Which three actions should you complete before you upgrade to a newer version of software? (Choose three.)

a) Review the release notes to determine any impact of upgrading to a newer version of software.

b) Ensure the firewall is connected to a reliable power source.

c) Export the device state.

d) Create and externally store a backup before you upgrade.

e) Put the firewall in maintenance mode.

Q4. After an Applications and Threats dynamic update is downloaded to the firewall, where can information about changes to the App-IDs be found?

a) Summary link in the log event detail reporting the dynamic update file download

b) Review Policies link at the bottom of the Security policy rules display

c) Review Apps link appearing next to the downloaded Applications and Threats file

d) Details link in the dynamic file availability announcement appearing in the News Feed widget on the dashboard

Q5. The GlobalProtect Data File dynamic update contains which kinds of data?

a) GlobalProtect client package software updates for Windows and Macintosh

b) list of available connection points for Prisma Access

c) HIP check detection data for the GlobalProtect clients

d) updates to cypher suites used by the GlobalProtect client

Q6. When application details are viewed in the App-ID database, which field indicates that a recent change to the application might affect your Security policy rules?

a) Name

b) Depends on

c) Previously Identified As

d) App-ID Enabled

2.6 Identify what a security zone is and how to use it

Q1. Which two default zones are included with the PAN-OS software? (Choose two.)

a) Interzone

b) Extrazone

c) Intrazone

d) Extranet

Q2. Which two zone types are valid? (Choose two.)

a) trusted

b) tap

c) virtual wire

d) untrusted

e) dmz

Q3. Which two statements about interfaces are correct? (Choose two.)

a) Interfaces must be configured before you can create a zone.

b) Interfaces do not have to be configured before you can create a zone.

c) An interface can belong to only one zone.

d) An interface can belong to multiple zones.

Q4. Which two interface types can belong in a Layer 3 zone? (Choose two.)

a) Loopback

b) Tap

c) Tunnel

d) Virtual Wire

Q5. What are used to control traffic through zones?

a) access lists

b) Security policy lists

c) Security policy rules

d) Access policy rules

2.7 Identify and configure firewall interfaces

Q1. For inbound inspection, which two actions can be done with a Tap interface? (Choose two.)

a) encrypt traffic

b) decrypt traffic

c) allow or block traffic

d) log traffic

Q2. Which two actions can be done with a Virtual Wire interface? (Choose two.)

a) NAT

b) route

c) switch

d) log traffic

Q3. Which two actions can be done with a Layer 3 interface? (Choose two.)

a) NAT

b) route

c) switch

d) create a Virtual Wire object

Q4. Layer 3 interfaces support which two items? (Choose two.)

a) NAT

b) IPv6

c) switching

d) spanning tree

Q5. Layer 3 interfaces support which three advanced settings? (Choose three.)

a) IPv4 addressing

b) IPv6 addressing

c) NDP configuration

d) link speed configuration

e) link duplex configuration

Q6. Layer 2 interfaces support which three items? (Choose three.)

a) spanning tree blocking

b) traffic examination

c) forwarding of spanning tree BPDUs

d) traffic shaping via QoS

e) firewall management

f) routing

Q7. Which two interface types support subinterfaces? (Choose two.)

a) Virtual Wire

b) Layer 2

c) Loopback

d) Tunnel

Q8. Which two statements are true regarding Layer 3 interfaces? (Choose two.)

a) You can configure a Layer 3 interface with one or more IP addresses as a DHCP client.

b) A Layer 3 interface can only have one DHCP assigned address.

c) You can assign only one IPv4 addresses to the same interface.

d) You can enable an interface to send IPv4 Router Advertisements by selecting the Enable Router Advertisement check box on the Router Advertisement tab.

e) You can apply an Interface Management Profile to the interface.

Q9. Which statement is true regarding aggregate Ethernet interfaces?

a) Members of an Aggregate Interface Group can be of different media types.

b) An Aggregate Interface Group can be set to a type of tap.

c) Member Ethernet interfaces of an Aggregate Interface Group must have the same transmission speeds.

d) A Layer 3 Aggregate Interface Group can have more than one IP assigned to it.

e) Member Ethernet interfaces can be assigned to different virtual routers.

2.8 Configure a virtual router

Q1. What is the default administrative distance of a static route within the PAN-OS software?

a) 1

b) 5

c) 10

d) 100

Q2. Which two dynamic routing protocols are available in the PAN-OS software? (Choose two.)

a) RIP1

b) RIPv2

c) OSPFv3

d) EIGRP

Q3. Which value is used to distinguish the preference of routing protocols?

a) metric

b) weight

c) distance

d) cost

e) administrative distance

Q4. Which value is used to distinguish the best route within the same routing protocol?

a) metric

b) weight

c) distance

d) cost

e) administrative distance

Q5. In path monitoring, what is used to monitor remote network devices?

a) ping

b) SSL

c) HTTP

d) HTTPS

e) link state

Domain 3 – Managing Objects

3.2 Identify how to create services

Q1. Which two statements are true about a Role Based Admin Role Profile role? (Choose two.)

a) It is a built-in role.

b) It can be used for CLI commands.

c) It can be used for XML API.

d) Superuser is an example.

Q2. The management console supports which two authentication types? (Choose two.)

a) RADIUS

b) SMB

c) LDAP

d) TACACS+

e) AWS

Q3. Which two Dynamic Admin Role types are available on the PAN-OS software? (Choose two.)

a) superuser

b) superuser (write only)

c) device user

d) device administrator (read-only)

Q4. Which type of profile does an Authentication Sequence include?

a) Security

b) Authorization

c) Admin

d) Authentication

Q5. An Authentication Profile includes which other type of profile?

a) Server

b) Admin

c) Customized

d) Built-In

Q6. True or false: Dynamic Admin Roles are called "dynamic" because you can customize them.

a) true

b) false

Q7. Which profile is used to override global Minimum Password Complexity Requirements?

a) Authentication

b) Local

c) User

d) Password

3.4 Configure application filters and application groups

Q1. What does an application filter enable an administrator to do?

a) manually categorize multiple service filters

b) dynamically categorize multiple service filters

c) dynamically categorize multiple applications

d) manually categorize multiple applications

Q2. Which two items can be added to an application group? (Choose two.)

a) application groups

b) application services

c) application filters

d) application categories

Domain 4 - Policy Evaluation and Management

4.1 Identify the appropriate application-based security policy

Q1. What are two application dependencies for adobe-connectnow-base? (Choose two.)

a) ssl

b) skype

c) rtmp

d) adobe-base

e) ssh

4.2 Identify the purpose of specific security rule types

Q1. What are the two default (predefined) Security policy rule types in PAN-OS software? (Choose two.)

a) Universal

b) Interzone

c) Intrazone

d) Extrazone

Q2. True or false. Because the first rule that matches the traffic is applied, the more specific rules must follow the more general ones.

a) true

b) false

Q3. Which statement is true?

a) For Intrazone traffic, traffic logging is enabled by default.

b) For Interzone traffic, traffic logging is enabled by default.

c) For Universal traffic, traffic logging is enabled by default.

d) For any rule type, traffic logging is enabled by default.

4.3 Identify and configure Security policy match conditions, actions, and logging options

Q1. What are the two default (predefined) Security policy rule types in PAN-OS software? (Choose two.)

a) Universal

b) Interzone

c) Intrazone

d) Extrazone

Q2. What will be the result of one or more occurrences of shadowing?

a) a failed commit

b) an invalid configuration

c) a warning

d) an alarm window

Q3. Which type of Security policy rules most often exist above the two predefined security policies?

a) Intrazone

b) Interzone

c) Universal

d) Global

Q4. What does the TCP Half Closed setting mean?

a) maximum length of time that a session remains in the session table between reception of the first FIN and reception of the third FIN or RST

b) minimum length of time that a session remains in the session table between reception of the first FIN and reception of the second FIN or RST

c) maximum length of time that a session remains in the session table between reception of the first FIN and reception of the second FIN or RST

d) minimum length of time that a session remains in the session table between reception of the first FIN and reception of the third FIN or RST.

Q5. What are two application characteristics? (Choose two.)

a) stateful

b) excessive bandwidth use

c) intensive

d) evasive

4.4 Identify and implement the proper NAT policy

Q1. What are two source NAT types? (Choose two.)

a) universal

b) static

c) dynamic

d) extrazone

Q2. Which phrase is a simple way to remember how to configure Security policy rules where NAT was implemented?

a) post-NAT IP, pre-NAT zone

b) post-NAT IP, post-NAT zone

c) pre-NAT IP, post-NAT zone

d) pre-NAT IP, pre-NAT zone

Q3. What are two types of destination NAT? (Choose two.)

a) dynamic IP (with session distribution)

b) DIPP

c) global

d) static

Q4. What are two possible values for DIPP NAT oversubscription? (Choose two.)

a) 1x

b) 4x

c) 16x

d) 32x

Q4. Which statement is true regarding bidirectional NAT?

a) For static translations, bidirectional NAT enables the firewall to create a corresponding translation in the opposite direction of the translation you configure.

b) For static translations, bidirectional NAT enables the firewall to create a corresponding translation in the same direction of the translation you configure.

c) For dynamic translations, bidirectional NAT enables the firewall to create a corresponding translation in the opposite direction of the translation you configure.

d) For dynamic translations, bidirectional NAT enables the firewall to create a corresponding translation in the same direction of the translation you configure.

4.5 Identify the tools available to optimize Security policies

Q1. The Policy Optimizer does not analyze which statistics?

a) applications allowed through port-based Security policy rules

b) the usage of existing App-IDs in Security policy rules

c) which users matched security policies

d) existing Security policy rule App-IDs that have not matched processed traffic

e) days since the latest new application discovery in a port-based Security policy rule

Q2. Which column in the **Applications and Threats** screen includes the options **Review Apps** and **Review Policies**?

a) Features

b) Type

c) Version

d) Action

Q3. Which link can you select in the web interface to minimize the risk using of installing new App-ID updates?

a) Enable new apps in content

b) Disable new apps in App-ID database

c) Disable new apps in content update

d) Enable new apps in App-ID database

Q4. Which two protocols are implicitly allowed when you select the facebook-base application? (Choose two.)

a) web-browsing

b) chat

c) gaming

d) ssl

Domain 5 – Securing Traffic

5.1 Identify and apply the appropriate Security Profile

Q1. What are two benefits of Vulnerability Protection Security Profiles? (Choose two.)

a) prevent compromised hosts from trying to communicate with external C2 servers

b) protect against viruses, worms, and Trojans

c) prevent exploitation of system flaws

d) prevent unauthorized access to systems

Q2. A URL Filtering Profile is part of which type of identification?

a) App-ID

b) Content-ID

c) User-ID

d) Service

Q3. Which stage of the attack lifecycle is most likely to be stopped by dividing the network into separate security zones?

a) Reconnaissance

b) Execution

c) Lateral movement

d) Data exfiltration

Q4. Which component can tell you if an attack is an APT or a broad attack designed to produce a botnet for future abuse?

a) next-generation firewall

b) WildFire

c) MineMeld

d) AutoFocus

Q5. Packet Buffer Protection defends against which type of denial-of-service attack?

a) from distributed sessions

b) from a single App-ID source

c) from multiple App-ID sources

d) from a single session

Q6. Which defense is turned on when a Packet Buffer Protection event is detected? (Choose two.)

a) SYN cookie management of attacking session traffic

b) Global Random Early Drop of packets from the attacking session

c) block all packets from the attacking session for the configured duration if the attack persists for a certain configured time

d) block all packets from the attacking IP address for the configured duration if the attack persists for a certain configured time

Q7. What are the two components of Denial-of-Service Protection? (Choose two.)

a) Zone Protection Profile

b) DoS Protection Profile and policy rules

c) load protection

d) reconnaissance protection

Q8. Which statement describes the new machine learning capabilities implemented within security profiles introduced in PAN-OS 10.0?

a) Machine learning can be performed by the firewall on the stream of data passing through it, allowing threats to be blocked without signatures.

b) Machine learnt models can be implemented by the firewall on the stream of data passing through it, allowing threats to be blocked without signatures.

c) Machine learnt models can be implemented by the firewall, but only to detect threats after they have passed through the firewall.

d) Machine learning can be performed by the firewall on the stream of data passing through it, identifying threats that have already passed through the firewall.

5.2 Identify the difference between Security policy actions and Security Profile actions

Q1. Which two actions are available for Antivirus Security Profiles? (Choose two.)

a) continue

b) allow

c) block IP

d) alert

Q2. Which two HTTP Header Logging options are within a URL Filtering Profile? (Choose two.)

a) User-Agent

b) Safe Search

c) URL redirection

d) X-Forwarded-For

5.3 Identify how the firewall can use the cloud DNS Security to control traffic based on domains

Q1. Which two actions are required to implement DNS Security inspections of traffic? (Choose two.)

a) add an Anti-Spyware Security Profile with DNS remediations to a Security policy

b) enabled the Advanced DNS Security check box in General Settings

c) configure an Anti-Spyware Security Profile with DNS remediations

d) enter the address for the Secure DNS Service in the firewalls DNS settings

5.4 Identify how the firewall can use the PAN-DB database to control traffic based on websites

Q1. Which two types of attacks does the PAN-DB prevent? (Choose two.)

a) phishing sites

b) HTTP-based command and control

c) infected JavaScript

d) flood attacks

5.5 Identify how to control access to specific URLs using custom URL filtering categories

Q1. Which two valid URLs can be used in a custom URL category? (Choose two.)

a) ww.youtube.**

b) www.**.com

c) www.youtube.com

d) *youtube*

e) *.youtube.com

Continuing Your Learning Journey with Palo Alto Networks

Training from Palo Alto Networks and our Authorized Training Partners delivers the knowledge and expertise to prepare you to protect our way of life in the digital age. Our trusted security certifications give you the Palo Alto Networks product portfolio knowledge necessary to prevent successful cyberattacks and to safely enable applications.

Digital Learning

For those of you who want to keep up to date on our technology, a learning library of *free* digital learning is available. These on-demand, self-paced digital learning classes are a helpful way to reinforce the key information for those who have been to the formal hands-on classes. They also serve as a useful overview and introduction to working with our technology for those unable to travel to a hands-on, instructor-led class.

Simply register in our Learning Center and you will be given access to our digital learning portfolio. These online classes cover foundational material and contain narrated slides, knowledge checks, and, where applicable, demos for you to access.

New courses are being added often, so check back to see new curriculum available.

Instructor-Led Training

Looking for a hands-on, instructor-led course in your area?

Palo Alto Networks Authorized Training Partners (ATPs) are located globally and offer a breadth of solutions from onsite training to public, open environment classes. About 42 authorized training centers are delivering online courses in 14 languages and at convenient times for most major markets worldwide. For class schedule, location, and training offerings, see https://www.paloaltonetworks.com/services/education/atc-locations.

Learning Through the Community

You also can learn from peers and other experts in the field. Check out our communities site at https://live.paloaltonetworks.com, where you can:

- Discover reference material
- Learn best practices
- Learn what is trending