DUMPSARENA

Palo Alto Networks Systems Engineer (PSE): Software Firewall Professional

Palo Alto Networks PSE-SoftwareFirewall

Version Demo

Total Demo Questions: 7

Total Premium Questions: 79

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Which two methods of Zero Trust implementation can benefit an organization? (Choose two.)

A. Boundaries are established.

B. Security automation is seamlessly integrated.

C. Compliance is validated.

D. Access controls are enforced.

ANSWER: B D

Explanation:

Security automation is seamlessly integrated. is a key Zero Trust implementation method because Zero Trust relies on continuous assessment rather than a one-time trust decision. Integrating security tools and automating telemetry collection, policy updates, detection, and response enables the organization to apply consistent controls at the speed required by dynamic users, applications, workloads, and threats. Automation also helps operational teams correlate signals from identity, endpoint, network, and cloud environments so access decisions can adapt as risk changes.

Access controls are enforced. is equally fundamental because Zero Trust requires explicit, least-privilege access to protected resources. Access must be evaluated and enforced using contextual signals such as user identity, device posture, application, location, and risk. Enforcing granular controls limits users and systems to only the resources and actions required for their authorized task, reducing exposure and containing the impact of compromised credentials or endpoints. Palo Alto Networks describes Zero Trust as a model centered on continuously verifying access and enforcing least privilege, while its Zero Trust guidance emphasizes automated security operations and policy enforcement. See Palo Alto Networks: What Is Zero Trust? and Zero Trust Network Security.

QUESTION NO: 2

Which of the following can provide application-level security for a web-server instance on Amazon Web Services (AWS)?

A. VM-Series firewalls

B. Hardware firewalls

C. Terraform templates

D. Security groups

ANSWER: A

Explanation:

VM-Series firewalls provide application-level security for web-server instances deployed in AWS. The VM-Series is Palo Alto Networks' virtualized next-generation firewall, designed to run in public-cloud environments including AWS. It inspects traffic using App-ID to identify and control applications regardless of port or protocol, rather than relying only on network-layer addressing and ports. For a web-server workload, this enables granular policies governing web applications and web traffic, while also applying security capabilities such as Threat Prevention, URL Filtering, WildFire analysis, and decryption where appropriate. The firewall can be deployed inline with application traffic, such as in a centralized inspection architecture or a workload-specific design, so that traffic to and from the web-server instance is subject to consistent next-generation security policy enforcement. This makes VM-Series an appropriate choice when the requirement specifically calls for application-aware visibility and protection in AWS. Palo Alto Networks documents AWS deployment architectures and VM-Series cloud capabilities at Set Up the VM-Series Firewall on AWS and VM-Series Virtual Firewalls.

QUESTION NO: 3

Why are VM-Series firewalls and hardware firewalls that are external to the Kubernetes cluster problematic for protecting containerized workloads?

A. They function differently based on whether they are located inside or outside of the cluster.

B. They are located outside the cluster and have no visibility into application-level cluster traffic.

C. They are managed by another entity when located inside the cluster.

D. They do not scale independently of the Kubernetes cluster.

ANSWER: B

Explanation:

They are located outside the cluster and have no visibility into application-level cluster traffic. Kubernetes workloads communicate extensively through dynamic, east-west traffic between pods and services within the cluster. This traffic can remain internal to the Kubernetes networking layer and may not traverse an external VM-Series or hardware firewall. As a result, an external firewall does not have the necessary context to consistently inspect and enforce policy for the individual application connections occurring between containerized workloads.

Effective Kubernetes security requires visibility that aligns with Kubernetes constructs such as namespaces, labels, pods, and services, while also accommodating rapid workload creation, deletion, and rescheduling. Palo Alto Networks positions the CN-Series firewall for this purpose: it is deployed within the Kubernetes environment and integrates with the cluster to apply application-aware security controls to container traffic. This placement provides the workload-level visibility needed to protect internal application communications, rather than limiting enforcement to traffic that enters or leaves the cluster. See the Palo Alto Networks CN-Series documentation and the Palo Alto Networks container security overview.

QUESTION NO: 4

Which two subscriptions should be recommended to a customer who is deploying VM-Series firewalls to a private data center but is concerned about protecting data-center resources from malware and lateral movement? (Choose two.)

A. Threat Prevention

B. SD-WAN

C. Intelligent Traffic Offload

D. WildFire

ANSWER: A D

Explanation:

Threat Prevention and **WildFire** directly address the stated need to protect private data-center workloads from malware and threats that can spread laterally between systems. Threat Prevention supplies the firewall's inline prevention capabilities, including Vulnerability Protection, Anti-Spyware, Antivirus, and DNS Security services. Applied to security policies governing east-west as well as north-south traffic, these profiles inspect allowed application traffic and block known exploits, command-and-control activity, malicious DNS activity, and malware. This gives VM-Series firewalls an essential control point for limiting compromise and movement within the data center.

WildFire complements that inline protection by analyzing unknown or suspicious files and producing malware verdicts and threat intelligence. WildFire analysis helps identify previously unseen and evasive malware that may target servers or propagate between data-center segments. The resulting protections are distributed through Palo Alto Networks content updates and can be enforced by security profiles, strengthening defenses against new threats in addition to known signatures. Together, these subscriptions provide prevention for known threats and cloud-delivered analysis for unknown malware, which is the appropriate combination for this deployment objective. See the Threat Prevention documentation and the WildFire documentation.

QUESTION NO: 5

Which two routing options are supported by VM-Series? (Choose two.)

A. RIP

B. OSPF

C. IGRP

D. BGP

ANSWER: B D

Explanation:

OSPF and **BGP** are supported dynamic-routing options for VM-Series firewalls. OSPF enables the firewall to exchange and calculate intra-domain routes using link-state routing, which is useful when the firewall participates in an enterprise routing domain. BGP enables policy-based route exchange with peers and is commonly used for cloud, data-center, WAN, and inter-domain designs. VM-Series can run these protocols through a virtual router, allowing dynamically learned routes to be used for forwarding decisions and advertised according to the configured routing policy. Palo Alto Networks documents BGP and OSPF as supported routing capabilities for VM-Series deployments, including cloud architectures in which dynamic route exchange is required for resilient connectivity. The relevant configuration concepts, such as virtual routers, routing-protocol profiles, route redistribution, and peer configuration, are managed in PAN-OS. See the Palo Alto Networks VM-Series Deployment Guide and the PAN-OS documentation for configuring BGP.

QUESTION NO: 6

Which type of group allows sharing cloud-learned tags with on-premises firewalls?

A. Notify •

B. Address

C. Template

D. Device

ANSWER: D

Explanation:

Device is correct because a Panorama device group is the policy and object-management scope used to associate managed firewalls with common configuration and policy data. For hybrid deployments, device groups provide the administrative boundary through which cloud-learned tag information can be made available to on-premises firewalls. The on-premises firewalls in the applicable device-group hierarchy can then use those tag values in policy constructs such as Dynamic Address Groups, allowing policy to follow changing cloud workload attributes instead of relying on manually maintained IP-address lists.

This approach is particularly useful when cloud resources scale, move, or are replaced frequently. Tags learned from cloud integrations represent workload context, and making them available through the device-group scope lets security teams apply consistent segmentation and access-control policy across cloud and on-premises environments. Panorama device groups are designed to centrally manage policy and objects for groups of firewalls, including inheritance through the device-group hierarchy. Palo Alto Networks documents device groups as the central mechanism for grouping firewalls that require common policy and object configuration, and documents Dynamic Address Groups as consumers of tag-based membership for policy enforcement. See Panorama Device Groups and Dynamic Address Groups.

QUESTION NO: 7

What is a design consideration for a prospect who wants to deploy VM-Series firewalls in an Amazon Web Services (AWS) environment?

A. Resources are shared within the cluster.

B. Only active-passive high availability (HA) is supported.

C. High availability (HA) clusters are limited to fewer than 8 virtual appliances.

D. Special AWS plugins are needed for load balancing.

ANSWER: B

Explanation:

Only active-passive high availability (HA) is supported. When designing a traditional VM-Series HA deployment in AWS, the firewall pair is deployed with one peer actively processing traffic and a passive peer ready to take over if the active firewall or its monitored resources fail. This model is necessary because AWS networking constructs do not provide the same Layer 2 adjacency and gratuitous ARP behavior available in many on-premises HA designs. Instead, the deployment uses AWS API integration and route, address, or interface updates to redirect traffic during failover. The design must therefore account for the AWS permissions, IAM role or credentials, subnet and routing layout, and monitoring configuration required for the passive firewall to assume traffic forwarding after a failure. Palo Alto Networks documents the AWS HA workflow as an active-passive configuration and provides the required cloud integration steps for failover. See the VM-Series High Availability for AWS documentation and the VM-Series on AWS deployment overview.