PCSFE Q&As

Palo Alto Networks Certified Software Firewall Engineer (PCSFE)

Pass Palo Alto Networks PCSFE Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/pcsfe.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Palo Alto Networks Official Exam Center

QUESTION 1

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

A. Special AWS plugins are needed for load balancing.

B. Resources are shared within the cluster.

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

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

Correct Answer: C

Explanation: A design consideration for a prospect who wants to deploy VM-Series firewalls in an Amazon Web Services (AWS) environment is that only active-passive high availability (HA) is supported. High availability (HA) is a feature that provides redundancy and failover protection for firewalls in case of hardware or software failure. Active-passive HA is a mode of HA that consists of two firewalls in a pair, where one firewall is active and handles all traffic, while the other firewall is passive and acts as a backup. Active-passive HA is the only mode of HA that is supported for VM-Series firewalls in an AWS environment, due to the limitations of AWS networking and routing. Active-active HA, which is another mode of HA that consists of two firewalls in a pair that both handle traffic and synchronize sessions, is not supported for VM-Series firewalls in an AWS environment. A design consideration for a prospect who wants to deploy VM-Series firewalls in an AWS environment is not that special AWS plugins are needed for load balancing, resources are shared within the cluster, or high availability (HA) clusters are limited to fewer than 8 virtual appliances, as those are not valid or relevant factors for firewall deployment in an AWS environment. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [High Availability Overview], [High Availability on AWS]

QUESTION 2

How are Palo Alto Networks Next-Generation Firewalls (NGFWs) deployed within a Cisco ACI architecture?

A. SDN code hooks can help detonate malicious file samples designed to detect virtual environments.

B. Traffic can be automatically redirected using static address objects.

C. Service graphs are configured to allow their deployment.

D. VXLAN or NVGRE traffic is terminated and inspected for translation to VLANs.

Correct Answer: C

Explanation: Palo Alto Networks Next-Generation Firewalls (NGFWs) are deployed within a Cisco ACI architecture using service graphs. Service graphs are logical representations of how traffic flows through different network services, such as firewalls, load balancers, or routers. By configuring service graphs, you can insert NGFWs into the traffic path and apply security policies to the traffic. References: [Palo Alto Networks NGFW Integration with Cisco ACI]

QUESTION 3

Which three NSX features can be pushed from Panorama in PAN-OS? (Choose three.)

A. Security group assignment of virtual machines (VMs)

B. Security groups

C. Steering rules

D. User IP mappings

E. Multiple authorization codes

Correct Answer: ABC

QUESTION 4

What is the structure of the YAML Ain\\'t Markup Language (YAML) file repository?

A. Deployment Type/Kubernetes/Environment

B. Kubernetes/Deployment Type/Environment

C. Kubernetes/Environment/Deplovment Type

D. Environment/Kubernetes/Deployment Type

Correct Answer: B

Explanation: Kubernetes/Deployment Type/Environment is the structure of the YAML Ain\\'t Markup Language (YAML) file repository. YAML is a human-readable data serialization language that is commonly used for configuration files. YAML file repository is a collection of YAML files that specify the resources and configuration for deploying and managing infrastructure components, such as firewalls, load balancers, networks, or servers. Kubernetes/Deployment Type/Environment is the structure of the YAML file repository that organizes the YAML files based on the following criteria: Kubernetes: The platform that provides orchestration, automation, and management of containerized applications. Deployment Type: The method or model of deploying and managing infrastructure components, such as Terraform, Ansible, Helm, or Kubernetes manifests. Environment: The type or stage of the cloud or virtualization environment, such as development, testing, staging, or production. Deployment Type/Kubernetes/Environment, Kubernetes/Environment/Deployment Type, and Environment/Kubernetes/Deployment Type are not the structure of the YAML file repository, but they are related ways of organizing YAML files based on different criteria. References: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [What is YAML?], [YAML File Repository]

QUESTION 5

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

A. Device

B. Notify

C. Address

D. Template

Correct Answer: C

Explanation: Address groups are the type of groups that allow sharing cloud-learned tags with on-premises firewalls. Address groups are dynamic objects that can include IP addresses or tags as members. Cloud-learned tags are tags that are assigned to cloud resources by cloud providers or third-party tools. By using address groups with cloud- learned tags, you can apply consistent security policies across your hybrid cloud environment. References: [Address Groups]

QUESTION 6

How is traffic directed to a Palo Alto Networks firewall integrated with Cisco ACI?

A. By using contracts between endpoint groups that send traffic to the firewall using a shared policy

B. Through a virtual machine (VM) monitor domain

C. Through a policy-based redirect

D. By creating an access policy

Correct Answer: C

Explanation: Traffic is directed to a Palo Alto Networks firewall integrated with Cisco ACI through a policy-based redirect. Cisco ACI is a software-defined network (SDN) solution that provides network automation, orchestration, and visibility. A policy-based redirect is a mechanism that allows Cisco ACI to redirect traffic from one endpoint group (EPG) to another EPG through a service device, such as a Palo Alto Networks firewall. The firewall can then inspect and enforce security policies on the redirected traffic before sending it back to Cisco ACI. Traffic is not directed to a Palo Alto Networks firewall integrated with Cisco ACI by using contracts between endpoint groups that send traffic to the firewall using a shared policy, through a virtual machine (VM) monitor domain, or by creating an access policy, as those are not valid methods for traffic redirection in Cisco ACI. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [Deploy the VM-Series Firewall on Cisco ACI], [Cisco ACI Policy-Based Redirect]

QUESTION 7

What is a benefit of network runtime security?

A. It more narrowly focuses on one security area and requires careful customization integration and maintenance

B. It removes vulnerabilities that have been baked into containers.

C. It is siloed to enhance workload security.

D. It identifies unknown vulnerabilities that cannot be identified by known Common Vulnerability and Exposure (CVE) lists.

Correct Answer: D

Explanation: A benefit of network runtime security is that it identifies unknown vulnerabilities that cannot be identified by known Common Vulnerability and Exposure (CVE) lists. Network runtime security is a type of security that monitors and analyzes network traffic in real time to detect and prevent malicious activities or anomalous behaviors. Network runtime security can identify unknown vulnerabilities that cannot be identified by known CVE lists, such as zero-day exploits, advanced persistent threats, or custom malware. Network runtime security can also provide visibility and context into network activity, such as application dependencies, user identities, device types, or threat intelligence. Network runtime security does not more narrowly focus on one security area and requires careful customization, integration, and maintenance, remove vulnerabilities that have been baked into containers, or is siloed to enhance workload security, as those are not benefits or characteristics of network runtime security. References: Palo Alto Networks Certified Software

Firewall Engineer (PCSFE), [Network Runtime Security], [What is CVE?]

QUESTION 8

Which two valid components are used in installation of a VM-Series firewall in an OpenStack environment? (Choose two.)

A. OpenStack heat template in JSON format

B. OpenStack heat template in YAML Ain\\'t Markup Language (YAML) format

C. VM-Series VHD image

D. VM-Series qcow2 image

Correct Answer: BD

Explanation: The two valid components that are used in installation of a VM-Series firewall in an OpenStack environment are: OpenStack heat template in YAML Ain\\'t Markup Language (YAML) format VM-Series qcow2 image OpenStack is a cloud computing platform that provides infrastructure as a service (IaaS) for deploying and managing virtual machines (VMs) and other resources. OpenStack environment requires network security that can protect the traffic between VMs or other cloud services from cyberattacks and enforce granular security policies based on application, user, content, and threat information. VM-Series firewall is a virtualized version of the Palo Alto Networks next-generation firewall that can be deployed on various cloud or virtualization platforms, including OpenStack. OpenStack heat template in YAML format is a valid component that is used in installation of a VM-Series firewall in an OpenStack environment. OpenStack heat template is a file that defines the resources and configuration for deploying and managing a VM-Series firewall instance on OpenStack. YAML is a human-readable data serialization language that is commonly used for configuration files. YAML format is supported for OpenStack heat templates for VM-Series firewalls. VM-Series qcow2 image is a valid component that is used in installation of a VM- Series firewall in an OpenStack environment. VM-Series qcow2 image is a file that contains the software image of the VM-Series firewall for OpenStack. qcow2 is a disk image format that supports features such as compression, encryption, snapshots, and copy-on-write. qcow2 format is supported for VM-Series images for OpenStack. OpenStack heat template in JSON format and VM-Series VHD image are not valid components that are used in installation of a VM-Series firewall in an OpenStack environment, as those are not supported formats for OpenStack heat templates or VM-Series images. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [Deploy the VM-Series Firewall on OpenStack], [What is YAML?], [What is qcow2?]

QUESTION 9

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

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

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

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

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

Correct Answer: A

Explanation: VM-Series firewalls and hardware firewalls that are external to the Kubernetes cluster are problematic for

protecting containerized workloads because they are located outside the cluster and have no visibility into application-level cluster traffic. Kubernetes is a platform that provides orchestration, automation, and management of containerized applications. Kubernetes cluster traffic consists of traffic between containers within a pod, across pods, or across namespaces. VM-Series firewalls and hardware firewalls that are external to the Kubernetes cluster cannot inspect or control this traffic, as they only see the encapsulated or aggregated traffic at the network layer. This creates blind spots and security gaps for containerized workloads. VM-Series firewalls and hardware firewalls that are external to the Kubernetes cluster are not problematic for protecting containerized workloads because they do not scale independently of the Kubernetes cluster, are managed by another entity when located inside the cluster, or function differently based on whether they are located inside or outside of the cluster, as those are not valid reasons or scenarios for firewall deployment in a Kubernetes environment. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [CN-Series Concepts], [VM-Series on Kubernetes]

QUESTION 10

What can be implemented in a CN-Series to protect communications between Dockers?

A. Firewalling

B. Runtime security

C. Vulnerability management

D. Data loss prevention (DLP)

Correct Answer: A

Explanation: CN-Series firewall can protect communications between Dockers by firewalling. Dockers are software platforms that provide containerization technology for packaging and running applications in isolated environments. Communications between Dockers are network connections between containers within a Docker host or across Docker hosts. CN-Series firewall is a containerized firewall that integrates with Kubernetes and provides visibility and control over container traffic. CN-Series firewall can protect communications between Dockers by firewalling, which is the process of inspecting and enforcing security policies on network traffic based on application, user, content, and threat information. CN-Series firewall can also leverage threat prevention technologies, such as antivirus, anti-spyware, vulnerability protection, URL filtering, file blocking, data filtering, and WildFire analysis, to block any malicious content or activity in the communications between Dockers. CN-Series firewall does not protect communications between Dockers by runtime security, vulnerability management, or data loss prevention (DLP), as those are not features or functions of CN-Series firewall. References: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [CN-Series Datasheet], [CN-Series Concepts], [What is Docker?]

QUESTION 11

Which two configuration options does Palo Alto Networks recommend for outbound high availability (HA) design in Amazon Web Services using a VM-Series firewall? (Choose two.)

A. Transit VPC and Security VPC

B. Traditional active-active HA

C. Transit gateway and Security VPC

D. Traditional active-passive HA

Correct Answer: CD

Explanation: Palo Alto Networks recommends two configuration options for outbound high availability (HA) design in Amazon Web Services using a VM-Series firewall: transit gateway and Security VPC, and traditional active-passive HA. Transit gateway and Security VPC allows you to use a single transit gateway to route traffic between multiple VPCs and the internet, while using a Security VPC to host the VM-Series firewalls. Traditional active-passive HA allows you to use two VM-Series firewalls in an HA pair, where one firewall is active and handles all traffic, while the other firewall is passive and takes over in case of a failure. References: [VM-Series Deployment Guide for AWS Outbound VPC]

QUESTION 12

What helps avoid split brain in active-passive high availability (HA) pair deployment?

A. Using a standard traffic interface as the HA2 backup

B. Enabling preemption on both firewalls in the HA pair

C. Using the management interface as the HA1 backup link

D. Using a standard traffic interface as the HA3 link

Correct Answer: C

Explanation: Using the management interface as the HA1 backup link helps avoid split brain in active-passive high availability (HA) pair deployment. High availability (HA) is a feature that provides redundancy and failover protection for firewalls in case of hardware or software failure. Active-passive HA is a mode of HA that consists of two firewalls in a pair, where one firewall is active and handles all traffic, while the other firewall is passive and acts as a backup. Split brain is a condition that occurs when both firewalls in an HA pair assume the active role and start processing traffic independently, resulting in traffic duplication, policy inconsistency, or session disruption. Split brain can be caused by network failures, device failures, or configuration errors that prevent the firewalls from communicating their HA status and synchronizing their configurations and sessions. Using the management interface as the HA1 backup link helps avoid split brain in active-passive HA pair deployment. The HA1 interface is used for exchanging HA state information and configuration synchronization between the firewalls. Using the management interface as the HA1 backup link provides redundancy and failover protection for the HA1 interface, ensuring that the firewalls can maintain their HA communication and avoid split brain. Using a standard traffic interface as the HA2 backup, enabling preemption on both firewalls in the HA pair, or using a standard traffic interface as the HA3 link do not help avoid split brain in active-passive HA pair deployment, but they are related features that can enhance performance and reliability. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [High Availability Overview], [Configure HA Backup Links], [Configure Heartbeat Backup]

QUESTION 13

Which two criteria are required to deploy VM-Series firewalls in high availability (HA)? (Choose two.)

A. Assignment of identical licenses and subscriptions

B. Deployment on a different host

C. Configuration of asymmetric routing

D. Deployment on same type of hypervisor

Correct Answer: AB

Explanation: To deploy VM-Series firewalls in high availability (HA), you need to assign identical licenses and subscriptions, and deploy them on a different host. Assigning identical licenses and subscriptions ensures that both firewalls have the same features and capabilities. Deploying them on a different host ensures that they are not affected by the same host failure. References: [VM-Series High Availability]

QUESTION 14

Which offering inspects encrypted outbound traffic?

A. WildFire

B. TLS decryption

C. Content-ID

D. Advanced URL Filtering (AURLF)

Correct Answer: B

Explanation: TLS decryption is the offering that inspects encrypted outbound traffic. TLS decryption is a feature that allows the firewall to decrypt and inspect outbound SSL/TLS traffic from internal clients to external servers. TLS decryption can inspect encrypted outbound traffic by applying threat prevention technologies, such as antivirus, anti-spyware, vulnerability protection, URL filtering, file blocking, data filtering, and WildFire analysis, to the decrypted traffic and blocking any malicious content or activity. WildFire, Content-ID, and Advanced URL Filtering (AURLF) are not offerings that inspect encrypted outbound traffic, but they are related solutions that can enhance security and visibility. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [TLS Decryption Overview], [Threat Prevention Datasheet]

QUESTION 15

Regarding network segmentation, which two steps are involved in the configuration of a default route to an internet router? (Choose two.)

A. Select the Static Routes tab, then click Add.

B. Select Network > Interfaces.

C. Select the Config tab. then select New Route from the Security Zone Route drop-down menu.

D. Select Network > Virtual Router, then select the default link to open the Virtual Router dialog.

Correct Answer: AD

Explanation: To configure a default route to an internet router, you need to select Network > Virtual Router, then select the default link to open the Virtual Router dialog. Then, select the Static Routes tab, then click Add. You can then specify the destination as 0.0.0.0/0 and the next hop as the IP address of the internet router1. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE)