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.geekcert.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

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

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 2

What Palo Alto Networks software firewall protects Amazon Web Services (AWS) deployments with network security delivered as a managed cloud service?

A. VM-Series

B. Cloud next-generation firewall

C. CN-Series

D. Ion-Series Ion-Series

Correct Answer: B

Explanation: Cloud next-generation firewall is the Palo Alto Networks software firewall that protects Amazon Web Services (AWS) deployments with network security delivered as a managed cloud service. Cloud next-generation firewall is a cloud-native solution that provides comprehensive security and visibility across AWS environments, including VPCs, regions, accounts, and workloads. Cloud next-generation firewall is deployed and managed by Palo Alto Networks as a service, eliminating the need for customers to provision, configure, or maintain any infrastructure or software. VM-Series, CN-Series, and Ion-Series are not Palo Alto Networks software firewalls that protect AWS deployments with network security delivered as a managed cloud service, but they are related solutions that can be

deployed on AWS or other platforms. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [Cloud Next-Generation Firewall Datasheet], [VM-Series Datasheet], [CN-Series Datasheet], [Ion-Series Datasheet]

QUESTION 3

A customer in a VMware ESXi environment wants to add a VM-Series firewall and partition an existing group of virtual machines (VMs) in the same subnet into two groups. One group requires no additional security, but the second group requires substantially more security.

How can this partition be accomplished without editing the IP addresses or the default gateways of any of the guest VMs?

A. Edit the IP address of all of the affected VMs. www*

B. Create a new virtual switch and use the VM-Series firewall to separate virtual switches using virtual wire mode. Then move the guests that require more security into the new virtual switch.

C. Create a Layer 3 interface in the same subnet as the VMs and then configure proxy Address Resolution Protocol (ARP).

D. Send the VLAN out of the virtual environment into a hardware Palo Alto Networks firewall in Layer 3 mode. Use the same IP address as the old default gateway, then delete it.

Correct Answer: B

Explanation: The partition can be accomplished without editing the IP addresses or the default gateways of any of the guest VMs by creating a new virtual switch and using the VM-Series firewall to separate virtual switches using virtual wire mode. Then move the guests that require more security into the new virtual switch. A virtual switch is a software-based switch that connects virtual machines (VMs) in a VMware ESXi environment. A virtual wire is a deployment mode of the VM-Series firewall that allows it to act as a bump in the wire between two network segments, without requiring an IP address or routing configuration. By creating a new virtual switch and using the VM-Series firewall to separate virtual switches using virtual wire mode, the customer can isolate the group of VMs that require more security from the rest of the network, and apply security policies to the traffic passing through the firewall. The partition cannot be accomplished without editing the IP addresses or the default gateways of any of the guest VMs by editing the IP address of all of the affected VMs, creating a Layer 3 interface in the same subnet as the VMs and then configuring proxy Address Resolution Protocol (ARP), or sending the VLAN out of the virtual environment into a hardware Palo Alto Networks firewall in Layer 3 mode. Use the same IP address as the old default gateway, then delete it, as those methods would require changing the network configuration of the guest VMs or introducing additional complexity and latency. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [Deploying Virtual Switches], [Virtual Wire Deployment], [Deploying Virtual Wire on VMware ESXi]

QUESTION 4

Which technology allows for granular control of east-west traffic in a software-defined network?

A. Routing

B. Microseqmentation

C. MAC Access Control List

D. Virtualization

Correct Answer: B

Explanation: Microsegmentation is a technology that allows for granular control of east- west traffic in a software-defined network. Microsegmentation divides the network into smaller segments or zones based on application or workload characteristics, and applies security policies to each segment. This reduces the attack surface and prevents unauthorized access or lateral movement within the network. Routing, MAC Access Control List, and Virtualization are not technologies that provide microsegmentation, but they are related concepts that can be used in conjunction with microsegmentation. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [Microsegmentation with Palo Alto Networks], [Microsegmentation for Dummies]

QUESTION 5

How are CN-Series firewalls licensed?

A. Data-plane vCPU

B. Service-plane vCPU

C. Management-plane vCPU

D. Control-plane vCPU

Correct Answer: A

Explanation: CN-Series firewalls are licensed by data-plane vCPU. Data-plane vCPU is the number of virtual CPUs assigned to the data plane of the CN-Series firewall instance. The data plane is the part of the CN-Series firewall that processes network traffic and applies security policies. CN-Series firewalls are licensed by data-plane vCPU, which determines the performance and capacity of the CN-Series firewall instance, such as throughput, sessions, policies, rules, and features. CN-Series firewalls are not licensed by service-plane vCPU, management-plane vCPU, or control-plane vCPU, as those are not factors that affect the licensing cost or consumption of CN-Series firewalls. References: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [CN-Series Licensing], [CN-Series System Requirements], [CN-Series Architecture]