Vendor: Palo Alto Networks

Exam Code:PCSFE

Exam Name: Palo Alto Networks Certified Software Firewall Engineer (PCSFE)

Version: Demo

QUESTION 1

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 2

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 3

Which Palo Alto Networks firewall provides network security when deploying a microservices-based application?

A. PA-Series

B. ICN-Series

C. VM-Series

D. HA-Series

Correct Answer: B

Explanation: CN-Series firewall is the Palo Alto Networks firewall that provides network security when deploying a microservices-based application. A microservices-based application is an application that consists of multiple independent and loosely coupled services that communicate with each other through APIs. A microservices-based application requires network security that can protect the inter-service communication from cyberattacks and enforce granular security policies based on application or workload characteristics. CN-Series firewall is a containerized firewall that integrates with Kubernetes and provides visibility and control over container traffic. CN-Series firewall can provide network security when deploying a microservices-based application by inspecting and enforcing security policies on traffic between containers within a pod, across pods, or across namespaces in a Kubernetes cluster. PA-Series, VM-Series, and HA-Series are not Palo Alto Networks firewalls that provide network security when deploying a microservices- based application, but they are related solutions that can be deployed on different platforms or environments. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [CN-Series Datasheet], [CN-Series Concepts], [What is a Microservices Architecture?]

QUESTION 4

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 5

How must a Palo Alto Networks Next-Generation Firewall (NGFW) be configured in order to secure traffic in a Cisco ACI environment?

A. It must be deployed as a member of a device cluster

B. It must use a Layer 3 underlay network

C. It must receive all forwarding lookups from the network controller

D. It must be identified as a default gateway

Correct Answer: B

Explanation: A Palo Alto Networks Next-Generation Firewall (NGFW) must be configured to use a Layer 3 underlay network in order to secure traffic in a Cisco ACI environment. A Layer 3 underlay network is a physical network that provides IP connectivity between devices, such as routers, switches, and firewalls. A Palo Alto Networks NGFW must use a Layer 3 underlay network to communicate with the Cisco ACI fabric and receive traffic redirection from the Cisco ACI policy-based redirect mechanism. A Palo Alto Networks NGFW does not need to be deployed as a member of a device cluster, receive all forwarding lookups from the network controller, or be identified as a default gateway in order to secure traffic in a Cisco ACI environment, as those are not valid requirements or options for firewall integration with Cisco ACI. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [Deploy the VM-Series Firewall on Cisco ACI], [Cisco ACI Underlay Network]

QUESTION 6

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 7

What is required to integrate a Palo Alto Networks VM-Series firewall with Azure Orchestration?

A. Aperture orchestration engine

B. Client-ID

C. Dynamic Address Groups

D. API Key

Correct Answer: D

Explanation: To integrate a Palo Alto Networks VM-Series firewall with Azure Orchestration, you need an API Key. The API Key is used to authenticate and authorize requests from Azure Orchestration to the VM-Series firewall. The API Key is generated on the VM-Series firewall and copied to Azure Orchestration. References: [Azure Orchestration Integration with Palo Alto Networks VM-Series Firewalls]

QUESTION 8

A CN-Series firewall can secure traffic between which elements?

A. Host containers

B. Source applications

C. Containers

D. IPods

Correct Answer: C

Explanation: Containers are the elements that a CN-Series firewall can secure traffic between. Containers are isolated units of software that run on a shared operating system and have their own resources, dependencies, and configuration. A CN-Series firewall can inspect and enforce security policies on traffic between containers within a pod, across pods, or across namespaces in a Kubernetes cluster. Host containers, source applications, and IPods are not valid elements that a CN-Series firewall can secure traffic between. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [CN-Series Concepts], [What is a Container?]

QUESTION 9

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 10

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 11

Which software firewall would assist a prospect who is interested in securing extensive DevOps deployments?

A. CN-Series

B. Ion-Series

C. Cloud next-generation firewall

D. VM-Series

Correct Answer: D

Explanation: VM-Series firewall is the software firewall that would assist a prospect who is interested in securing extensive DevOps deployments. DevOps is a set of practices that combines software development and IT operations to deliver software products faster and more reliably. DevOps deployments require network security that can protect the traffic between different stages of the software development lifecycle, such as development, testing, staging, and production, as well as between different cloud or virtualization platforms, such as public clouds, private clouds, or on-premises data centers. 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. VM-Series firewall can assist a prospect who is interested in securing extensive DevOps deployments by providing comprehensive security and visibility across hybrid and multi-cloud environments, protecting applications and data from cyberattacks, and supporting automation and orchestration tools that simplify and accelerate the deployment and configuration of firewalls across different platforms. CN-Series, Ion-Series, and Cloud next-generation firewall are not software firewalls that would assist a prospect who is interested in securing extensive DevOps deployments, but they are related solutions that can be deployed on specific platforms or environments. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [VM-Series Datasheet], [VM-Series Deployment Guide], [What is DevOps?]

QUESTION 12

Which two factors lead to improved return on investment for prospects interested in Palo Alto Networks virtualized next-generation firewalls (NGFWs)? (Choose two.)

A. Decreased likelihood of data breach

B. Reduced operational expenditures

C. Reduced time to deploy

D. Reduced insurance premiums

Correct Answer: AC

Explanation: The two factors that lead to improved return on investment for prospects interested in Palo Alto Networks virtualized next-generation firewalls (NGFWs) are: Decreased likelihood of data breach Reduced time to deploy Palo Alto Networks virtualized NGFWs are virtualized versions of the Palo Alto Networks next-generation firewall that can be deployed on various cloud or virtualization platforms. Palo Alto Networks virtualized NGFWs provide comprehensive security and visibility across hybrid and multi-cloud environments, protecting applications and data from cyberattacks. By using Palo Alto Networks virtualized NGFWs, prospects can decrease the likelihood of data breach by applying granular security policies based on application, user, content, and threat information, and by leveraging cloud-delivered services such as Threat Prevention, WildFire, URL Filtering, DNS Security, and Cortex Data Lake. By using Palo Alto Networks virtualized NGFWs, prospects can also reduce the time to deploy by taking advantage of automation and orchestration tools such as Terraform, Ansible, CloudFormation, ARM templates, and Panorama plugins that simplify and accelerate the deployment and configuration of firewalls across different cloud platforms. Reduced operational expenditures and reduced insurance premiums are not factors that lead to improved return on investment for prospects interested in Palo Alto Networks virtualized NGFWs, but they may be potential benefits or outcomes of using them. References: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [VM-Series Datasheet], [CN-Series Datasheet], [Cloud Security Solutions]