Number: PCSFE
Passing Score: 800
Time Limit: 120
File Version: 14.0

Exam Code: PCSFE

Exam Name: Palo Alto Networks Certified Software Firewall Engineer

Exam A

QUESTION 1
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
Section:
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.
Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [CNSeries Datasheet], [CN-Series Concepts], [What is Docker?]

QUESTION 2
Which two public cloud platforms does the VM-Series plugin support? (Choose two.)

A. Azure
B. lIBM Cloud
C. Amazon Web Services
D. lOCI

Correct Answer: A, C
Section:
Explanation:
The two public cloud platforms that the VM-Series plugin supports are:
Azure
Amazon Web Services (AWS)
A public cloud platform is a cloud computing service that provides infrastructure as a service (IaaS), platform as a service (PaaS), or software as a service (SaaS) to customers over the internet. A public cloud platform requires network security that can protect the traffic between different cloud services or regions 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. VM-Series plugin is a software component that extends the functionality of the VM-Series firewall and Panorama to support specific features and capabilities of different cloud platforms. Azure and AWS are two public cloud platforms that the VM-Series plugin supports. Azure is a public cloud platform that provides a range of cloud services, such as compute, storage, networking, databases, analytics, artificial intelligence, and more. AWS is a public cloud platform that provides a range of cloud services, such as EC2, S3, VPC, Lambda, and more. The VM-Series plugin supports Azure and AWS by enabling features such as bootstrapping, dynamic address groups, scaling, load balancing, high availability, monitoring, logging, and automation for VM-Series firewalls and Panorama on these platforms. IBM Cloud and OCI are not public cloud platforms that the VM-Series plugin supports, but they are related platforms that can be used for other purposes. Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [VM-Series Plugin Overview], [VM-Series Plugin for Azure], [VM-Series Plugin for AWS], [What is Azure?], [What is AWS?]

QUESTION 3
With which two private cloud environments does Palo Alto Networks have deep integrations?
(Choose two.)

A. VMware NSX-T
B. Cisco ACI
C. Dell APEX
D. Nutanix

Correct Answer: A, B
Section:
Explanation:
The two private cloud environments that Palo Alto Networks have deep integrations with are:
VMware NSX-T
Cisco ACI
A private cloud environment is a cloud computing service that provides infrastructure as a service (IaaS) or platform as a service (PaaS) to customers within a private network or data center. A private cloud environment requires network security that can protect the traffic between different virtual machines (VMs) or other resources from cyberattacks and enforce granular security policies based on application, user, content, and threat information. Palo Alto Networks have deep integrations with VMware NSX-T and Cisco ACI, which are two private cloud environments that provide network virtualization, automation, and security for cloud-native applications. VMware NSX-T is a private cloud environment that provides software-defined networking (SDN) and security for heterogeneous endpoints and workloads across multiple hypervisors, containers, bare metal servers, or clouds. Cisco ACI is a private cloud environment that provides application-centric infrastructure (ACI) and security for physical and virtual endpoints across multiple data centers or clouds. Palo Alto Networks have deep integrations with VMware NSX-T and Cisco ACI by enabling features such as dynamic address groups, service insertion, policy redirection, service chaining, orchestration, monitoring, logging, and automation for VM-Series firewalls and Panorama on these platforms. Dell APEX and Nutanix are not private cloud environments that Palo Alto Networks have deep integrations with, but they are related platforms that can be used for other purposes. Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [Deploy the VM-Series Firewall on VMware NSX-T], [Deploy the VM-Series Firewall on Cisco ACI], [What is VMware NSX-T?], [What is Cisco ACI?]

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
Section:
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. Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [What is YAML?], [YAML File Repository]

QUESTION 5
Which feature must be configured in an NSX environment to ensure proper operation of a VM-Series firewall in order to secure east-west traffic?

A. Deployment of the NSX DFW
B. VMware Information Sources
C. User-ID agent on a Windows domain server
D. Device groups within VMware Services Manager

Correct Answer: A
Section:
Explanation:
Deployment of the NSX Distributed Firewall (DFW) must be configured in an NSX environment to ensure proper operation of a VM-Series firewall in order to secure east-west traffic. East-west traffic is the traffic that flows between applications or workloads within a network or a cloud environment.
NSX environment is a private cloud environment that provides software-defined networking (SDN) and security for heterogeneous endpoints and workloads across multiple hypervisors, containers, bare metal servers, or clouds. NSX DFW is a feature that provides distributed stateful firewalling at the hypervisor level for every virtual machine (VM) in an NSX environment. Deployment of the NSX DFW must be configured in an NSX environment to ensure proper operation of a VM-Series firewall in order to secure east-west traffic by enabling features such as service insertion, policy redirection, service chaining, orchestration, monitoring, logging, and automation for VM-Series firewalls and Panorama on NSX environment. VMware Information Sources, User-ID agent on a Windows domain server, and device groups within VMware Services Manager do not need to be configured in an NSX environment to ensure proper operation of a VM-Series firewall in order to secure east-west traffic, as those are not required or relevant components for NSX integration. Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [Deploy the VM-Series Firewall on VMware NSX-T], [What is VMware NSX-T?], [What is NSX Distributed Firewall?]

QUESTION 6
Which two routing options are supported by VM-Series? (Choose two.)

A. OSPF
B. RIP
C. BGP
D. IGRP

Correct Answer: A, C
Section:
Explanation:
The two routing options that are supported by VM-Series are:
OSPF
BGP
Routing is a process that determines the best path for sending network packets from a source to a destination. Routing options are protocols or methods that enable routing between different networks or devices. VM-Series firewall is a virtualized version of the Palo Alto Networks nextgeneration firewall that can be deployed on various cloud or virtualization platforms. VM-Series firewall supports various routing options that allow it to participate in dynamic routing environments and exchange routing information with other routers or devices. OSPF and BGP are two routing options that are supported by VM-Series. OSPF is a routing option that uses link-state routing algorithm to determine the shortest path between routers within an autonomous system (AS). BGP is a routing option that uses path vector routing algorithm to determine the best path between routers across different autonomous systems (ASes). RIP and IGRP are not routing options that are supported by VM-Series, but they are related protocols that can be used for other purposes. Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [VM-Series Deployment Guide], [Routing Overview], [What is OSPF?], [What is BGP?]

QUESTION 7
What are two requirements for automating service deployment of a VM-Series firewall from an NSX Manager? (Choose two.)

A. vCenter has been given Palo Alto Networks subscription licenses for VM-Series firewalls.
B. Panorama has been configured to recognize both the NSX Manager and vCenter.
C. The deployed VM-Series firewall can establish communications with Panorama.
D. Panorama can establish communications to the public Palo Alto Networks update servers.

Correct Answer: B, C
Section:
Explanation:
The two requirements for automating service deployment of a VM-Series firewall from an NSX Manager are:
Panorama has been configured to recognize both the NSX Manager and vCenter.
The deployed VM-Series firewall can establish communications with Panorama.
NSX Manager is a software component that provides centralized management and control of the NSX environment, including network virtualization, automation, and security. Service deployment is a process that involves deploying and configuring network services, such as firewalls, load balancers, or routers, on the NSX environment. 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 NSX. Panorama is a centralized management server that provides visibility and control over multiple Palo Alto Networks firewalls and devices. Panorama has been configured to recognize both the NSX Manager and vCenter is a requirement for automating service deployment of a VM-Series firewall from an NSX Manager. vCenter is a software component that provides centralized management and control of the VMware environment, including hypervisors, virtual machines, and other resources. Panorama has been configured to recognize both the NSX Manager and vCenter by adding them as VMware service managers and enabling service insertion for VM-Series firewalls on NSX. This allows Panorama to communicate with the NSX Manager and vCenter, retrieve information about the NSX environment, and deploy and manage VM-Series firewalls as network services on the NSX environment. The deployed VM-Series firewall can establish communications with Panorama is a requirement for automating service deployment of a VM-Series firewall from an NSX Manager. The deployed VM-Series firewall can establish communications with Panorama by registering with Panorama using its serial number or IP address, and receiving configuration updates and policy rules from Panorama. This allows the VM-Series firewall to operate as part of the Panorama management domain, synchronize its settings and status with Panorama, and report its logs and statistics to Panorama. vCenter has been given Palo Alto Networks subscription licenses for VM-Series firewalls and Panorama can establish communications to the public Palo Alto Networks update servers are not requirements for automating service deployment of a VM-Series firewall from an NSX Manager, as those are not related or relevant factors for service deployment automation. Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [Deploy the VM-Series Firewall on VMware NSX-T], [Panorama Overview], [VMware Service Manager], [Register the Firewall with Panorama]

QUESTION 8
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
Section:
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 CNSeries 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. Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [CN-Series Licensing], [CN-Series System Requirements], [CN-Series Architecture]

QUESTION 9
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: A, D
Section:
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. Reference: Palo Alto Networks Certified Software Firewall Engineer (PCSFE)

QUESTION 10
Why are containers uniquely suitable for runtime security based on allow lists?

A. Containers have only a few defined processes that should ever be executed.
B. Developers define the processes used in containers within the Dockerfile.
C. Docker has a built-in runtime analysis capability to aid in allow listing.
D. Operations teams know which processes are used within a container.

Correct Answer: A
Section:
Explanation:
Containers are uniquely suitable for runtime security based on allow lists because containers have only a few defined processes that should ever be executed. Developers can specify the processes that are allowed to run in a container using a Dockerfile, but this does not guarantee that only those processes will run at runtime. Therefore, using an allow list approach can prevent any unauthorized or malicious processes from running in a container2. Reference: Container Security

QUESTION 11
Which two steps are involved in deployment of a VM-Series firewall on NSX? (Choose two.)

A. Create a virtual data center (vDC) and a vApp that includes the VM-Series firewall.
B. Obtain the Amazon Machine Images (AMIs) from marketplace.
C. Enable communication between Panorama and the NSX Manager.
D. Register the VM-Series firewall as a service.

Correct Answer: C, D
Section:
Explanation:
To deploy a VM-Series firewall on NSX, you need to enable communication between Panorama and the NSX Manager. This allows Panorama to receive information about the virtual machines and services in the NSX environment. You also need to register the VM-Series firewall as a service on the NSX Manager. This allows NSX to redirect traffic to the VM-Series firewall for inspection3.
Reference: VM-Series Deployment Guide for VMware NSX

QUESTION 12
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
Section:
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. Reference: [Palo Alto Networks NGFW

Integration with Cisco ACI]

QUESTION 13
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
Section:
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 VMSeries firewall. The API Key is generated on the VM-Series firewall and copied to Azure Orchestration.
Reference: [Azure Orchestration Integration with Palo Alto Networks VM-Series Firewalls]

QUESTION 14
Which service, when enabled, provides inbound traffic protection?

A. Advanced URL Filtering (AURLF)
B. Threat Prevention
C. Data loss prevention (DLP)
D. DNS Security

Correct Answer: D
Section:
Explanation:
DNS Security is a service that provides inbound traffic protection by preventing DNS-based attacks. DNS Security uses machine learning and threat intelligence to identify and block malicious domains, command and control (C2) traffic, and DNS tunneling. Reference: [DNS Security]

QUESTION 15
Which two features of CN-Series firewalls protect east-west traffic between pods in different trust zones? (Choose two.)

A. Intrusion prevention system
B. Communication with Panorama
C. External load balancer
D. Layer 7 visibility

Correct Answer: A, D
Section:
Explanation:
The two features of CN-Series firewalls that protect east-west traffic between pods in different trust zones are:
Intrusion prevention system
Layer 7 visibility
East-west traffic is the traffic that flows between applications or workloads within a network or a cloud environment. Pods are the smallest units of deployment in Kubernetes, consisting of one or more containers

that share resources and network space. Trust zones are segments of the network or the cloud environment that have different levels of security requirements or policies based on data sensitivity, user identity, device type, or application function. CN-Series firewalls are containerized firewalls that integrate with Kubernetes and provide visibility and control over container traffic.
Intrusion prevention system is a feature of CN-Series firewalls that protects east-west traffic between pods in different trust zones by detecting and blocking known exploits and vulnerabilities using signature-based and behavior-based methods. Layer 7 visibility is a feature of CN-Series firewalls that protects east-west traffic between pods in different trust zones by identifying and classifying applications and protocols based on their content and characteristics, regardless of port, encryption, or evasion techniques. Communication with Panorama and external load balancer are not features of CN-Series firewalls that protect east-west traffic between pods in different trust zones, but they are related features that can enhance management and performance. Reference: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [CN-Series Concepts], [CN-Series Deployment Guide for Native K8], [Intrusion Prevention System Overview], [App-ID Overview]

QUESTION 16
Which component can provide application-based segmentation and prevent lateral threat movement?

A. DNS Security
B. NAT
C. URL Filtering
D. App-ID

Correct Answer: D
Section:
Explanation:
App-ID is the component that can provide application-based segmentation and prevent lateral threat movement. Application-based segmentation is a method of dividing the network into smaller segments or zones based on application or workload characteristics, such as function, dependency, owner, or security posture. Lateral threat movement is a technique used by attackers to move across the network from one compromised host to another, looking for sensitive data or assets. App-ID is a feature that identifies and classifies applications and protocols based on their content and characteristics, regardless of port, encryption, or evasion techniques. App-ID can provide applicationbased segmentation and prevent lateral threat movement by applying granular security policies based on application information to each segment or connection, blocking unauthorized access or data exfiltration. DNS Security, NAT, and URL Filtering are not components that can provide application-based segmentation and prevent lateral threat movement, but they are related features that can enhance security and visibility. Reference: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [App-ID Overview], [Microsegmentation with Palo Alto Networks], [Lateral Movement]

QUESTION 17
What does the number of required flex credits for a VM-Series firewall depend on?

A. vCPU allocation
B. IP address allocation
C. Network interface allocation
D. Memory allocation

Correct Answer: A
Section:
Explanation:
The number of required flex credits for a VM-Series firewall depends on vCPU allocation. Flex credits are a flexible licensing model that allows customers to purchase and consume software NGFWs as needed, without having to specify the platform or deployment model upfront. Customers can use flex credits to provision VM-Series firewalls on any supported cloud or virtualization platform. The number of required flex credits for a VM-Series firewall depends on vCPU allocation, which is the number of virtual CPUs assigned to the VM-Series firewall instance. The vCPU allocation determines the performance and capacity of the VM-Series firewall instance, such as throughput, sessions, policies, rules, and features. The number of required flex credits for a VM-Series firewall does not depend on IP address allocation, network interface allocation, or memory allocation, as those are not factors that affect the licensing cost or consumption of flex credits. Reference: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [Flex Credits Datasheet], [Flex Credits FAQ], [VM-Series System Requirements]

QUESTION 18
Which element protects and hides an internal network in an outbound flow?

A. DNS sinkholing
B. User-ID
C. App-ID
D. NAT

Correct Answer: D
Section:
Explanation:
NAT is the element that protects and hides an internal network in an outbound flow. NAT is a feature that translates the source or destination IP address or port of a packet as it passes through the firewall. NAT can protect and hide an internal network in an outbound flow by replacing the private IP addresses of the internal hosts with a public IP address of the firewall or another device, making them appear as a single entity to the external network. This prevents external hosts from directly accessing or identifying the internal hosts, and also conserves the public IP address space. DNS sinkholing, User-ID, and App-ID are not elements that protect and hide an internal network in an outbound flow, but they are related features that can enhance security and visibility. Reference: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [NAT Overview], [DNS Sinkholing Overview], [User-ID Overview], [App-ID Overview]

QUESTION 19
How does a CN-Series firewall prevent exfiltration?

A. It employs custom-built signatures based on hash
B. It distributes incoming virtual private cloud (VPC) traffic across the pool of VM-Series firewalls.
C. It provides a license deactivation API key.
D. It inspects outbound traffic content and blocks suspicious activity.

Correct Answer: D
Section:
Explanation:
CN-Series firewall prevents exfiltration by inspecting outbound traffic content and blocking suspicious activity. Exfiltration is a technique used by attackers to steal sensitive data or assets from a compromised network or system, usually by sending them to an external destination, such as a command and control server, a drop zone, or an email address. CN-Series firewall is a containerized firewall that integrates with Kubernetes and provides visibility and control over container traffic. CNSeries firewall prevents exfiltration by inspecting outbound traffic content and blocking suspicious activity using threat prevention technologies, such as antivirus, anti-spyware, vulnerability protection, URL filtering, file blocking, data filtering, and WildFire analysis. CN-Series firewall does not prevent exfiltration by employing custom-built signatures based on hash, distributing incoming virtual private cloud (VPC) traffic across the pool of VM-Series firewalls, or providing a license deactivation API key, as those are not valid or relevant methods for exfiltration prevention.
Reference: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [CN-Series Concepts], [CN-Series Deployment Guide for Native K8], [Threat Prevention Datasheet], [What is Exfiltration?]

QUESTION 20
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
Section:
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. Reference: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [High Availability Overview], [Configure HA Backup Links], [Configure Heartbeat Backup]

QUESTION 21
What must be enabled when using Terraform templates with a Cloud next-generation firewall (NGFW) for Amazon Web Services (AWS)?

A. AWS CloudWatch logging
B. Access to the Cloud NGFW for AWS console
C. Access to the Palo Alto Networks Customer Support Portal
D. AWS Firewall Manager console access

Correct Answer: B
Section:
Explanation:
Access to the Cloud NGFW for AWS console must be enabled when using Terraform templates with a Cloud next-generation firewall (NGFW) for Amazon Web Services (AWS). Terraform is an open-source tool that allows users to define and provision infrastructure as code using declarative configuration files. Terraform templates are files that specify the resources and configuration for deploying and managing infrastructure components, such as firewalls, load balancers, networks, or servers. Cloud NGFW for AWS is a cloud-native solution that provides comprehensive security and visibility across AWS environments, including VPCs, regions, accounts, and workloads. Cloud NGFW for AWS 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. Access to the Cloud NGFW for AWS console must be enabled when using Terraform templates with a Cloud NGFW for AWS, as the console is the web-based interface that allows customers to view and manage their Cloud NGFW for AWS instances, policies, logs, alerts, and reports. The console also provides the necessary information and credentials for integrating with Terraform, such as the API endpoint, access key ID, secret access key, and customer ID. AWS CloudWatch logging, access to the Palo Alto Networks Customer Support Portal, and AWS Firewall Manager console access do not need to be enabled when using Terraform templates with a Cloud NGFW for AWS, as those are not required or relevant components for Terraform integration. Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [Terraform Overview], [Cloud Next-Generation Firewall Datasheet], [Cloud Next-Generation Firewall Deployment Guide], [Cloud Next-Generation Firewall Console Guide]

QUESTION 22
How does Prisma Cloud Compute offer workload security at runtime?

A. It automatically builds an allow-list security model for every container and service.
B. It quarantines containers that demonstrate increased CPU and memory usage.
C. It automatically patches vulnerabilities and compliance issues for every container and service.
D. It works with the identity provider (IdP; to identify overprivileged containers and services and it restricts network access

Correct Answer: A
Section:
Explanation:
Prisma Cloud Compute offers workload security at runtime by automatically building an allow-list security model for every container and service. Workload security is a type of security that protects applications and data from cyberattacks across different stages of the software development lifecycle, such as development, testing, staging, and production. Runtime security is a type of security that monitors and analyzes workload behavior in real time to detect and prevent malicious activities or anomalous behaviors. Prisma Cloud Compute is a cloud-native solution that provides comprehensive security and visibility across hybrid and multi-cloud environments, covering hosts, containers, serverless functions, and web applications. Prisma Cloud Compute offers workload security at runtime by automatically building an allow-list security model for every container and service, which defines the expected network connections, processes, file system activity, and system calls for each workload based on its baseline behavior. Prisma Cloud Compute

then enforces the allow-list security model and blocks any deviations or violations from the expected behavior. Prisma Cloud Compute does not quarantine containers that demonstrate increased CPU and memory usage, automatically patch vulnerabilities and compliance issues for every container and service, or work with the identity provider (IdP) to identify overprivileged containers and services and restrict network access, as those are not methods or features of Prisma Cloud Compute for workload security at runtime. Reference: [Palo Alto Networks Certified Software Firewall Engineer (PCSFE)], [Prisma Cloud Compute Datasheet], [Prisma Cloud Compute Overview], [Prisma Cloud Compute Runtime Defense]

QUESTION 23
Which software firewall would help a prospect interested in securing an environment with Kubernetes?

A. KN-Series
B. ML-Series
C. VM-Series
D. CN-Series

Correct Answer: D
Section:
Explanation:
CN-Series firewall is the software firewall that would help a prospect interested in securing an environment with Kubernetes. Kubernetes is a platform that provides orchestration, automation, and management of containerized applications. Kubernetes environment 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 help a prospect interested in securing an environment with Kubernetes by inspecting and enforcing security policies on traffic between containers within a pod, across pods, or across namespaces in a Kubernetes cluster. KN-Series, ML-Series, VM-Series, and Cloud next-generation firewall are not software firewalls that would help a prospect interested in securing an environment with Kubernetes, but they are related solutions that can be deployed on different platforms or environments. Reference: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [CNSeries Datasheet], [CN-Series Concepts], [What is Kubernetes?]

QUESTION 24
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: B, D
Section:
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. Reference: Palo Alto Networks Certified Software Firewall Engineer (PCSFE), [Deploy the VM-Series Firewall on OpenStack], [What is YAML?], [What is qcow2?]