DUMPSARENA

Palo Alto Networks System Engineer Professional - Strata Data Center

Palo Alto Networks PSE-StrataDC

Version Demo

Total Demo Questions: 8

Total Premium Questions: 87

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Center Architecture	39
Topic 2, Data Center Deployment	39
Topic 3, Data Center Operations	8
Topic 4, Data Center Troubleshooting	1
Total	87

QUESTION NO: 1

A company allows employees some personal use of the internet during work time However the CEO is concerned that employees are using too much of the bandwidth for YouTube. thus causing a performance problem. Which section of the SLR could confirm or allay this concern?

A. Categories with the Most Applications

B. Categories Consuming the Most Bandwidth

C. Bandwidth Consumed by Applications

D. High-Risk Applications

ANSWER: C

Explanation:

Bandwidth Consumed by Applications is the appropriate SLR section because the concern is specifically whether YouTube, an individual application, is responsible for a significant portion of network bandwidth use. This view attributes traffic volume to identified applications, allowing the engineer to find YouTube in the application results and assess its relative bandwidth consumption. The result can directly substantiate the CEO's concern when YouTube is a leading consumer, or allay it when its consumption is comparatively low and another application accounts for the traffic volume.

This application-level visibility is important because Palo Alto Networks identifies traffic by application rather than relying only on ports or broad web classifications. It supports a focused discussion of business impact and provides the evidence needed to consider an appropriate policy response, such as applying an application-specific QoS rule, limiting access during business hours, or continuing to permit reasonable personal use. Palo Alto Networks documents how application usage and bandwidth can be investigated through the Application Command Center and traffic monitoring in its Application Command Center documentation and logs and reports documentation.

QUESTION NO: 2

In a Cisco ACI deployment, which two constructs are used to insert a Palo Alto Networks firewall service into traffic between endpoint groups? (Choose two.)

A. Contract

B. Service graph

C. Bridge domain

D. Tenant

E. Virtual machine monitor domain

ANSWER: A B

Explanation:

A **contract** and a **service graph** are used for firewall service insertion in Cisco ACI. A contract defines the policy relationship between consumer and provider endpoint groups (EPGs). The service graph associates a Layer 4 through Layer 7 service, such as a Palo Alto Networks firewall, with that policy relationship so the ACI fabric can steer applicable traffic through the firewall.

An EPG identifies endpoints that share common policy requirements, but it does not independently insert a firewall service. A bridge domain provides Layer 2 forwarding context, and a tenant provides administrative isolation; neither replaces the service graph and contract used for application-policy-based firewall insertion. See Cisco's ACI Layer 4 through Layer 7 Services Configuration Guide and the Palo Alto Networks VM-Series Cisco ACI deployment guidance.

QUESTION NO: 3

Whichthree deployment modes of VM-Series firewalls are supported across NSX-T?

(Choose three)

A. Partner Service

B. Boot Strap

C. Prism Central

D. Tier-1 insertion

E. Tier-0 insertion

ANSWER: A D E

Explanation:

VM-Series firewalls integrate with VMware NSX-T as a partner service. This integration lets NSX-T Manager use the service definition and deployment details configured through Panorama to deploy and steer traffic through VM-Series firewall instances. Therefore, **Partner Service** is a supported NSX-T deployment model.

The two supported traffic-insertion deployment modes are **Tier-0 insertion** and **Tier-1 insertion**. Tier-0 insertion attaches the firewall service to a Tier-0 logical router, protecting traffic between the physical network and the NSX-T logical network. The firewall is deployed in virtual wire mode, and NSX-T redirects traffic crossing the logical router to the firewall for inspection. Tier-1 insertion similarly attaches the firewall service to a Tier-1 logical router, allowing inspection of traffic associated with segments and their uplink connectivity toward Tier-0 routers.

In both insertion models, Panorama centrally manages the VM-Series firewalls and their security policy, while NSX-T controls service insertion and traffic redirection. Palo Alto Networks documents these supported deployments at Supported Deployments of the VM-Series Firewall on VMware NSX-T and provides NSX-T deployment guidance in the VM-Series Documentation.

QUESTION NO: 4

Which environment is least likely to be placed on a public cloud by a hospital that has a large health information management application?

A. production

B. development

C. testing

D. QA

ANSWER: A

Explanation:

Production is the environment least likely to be placed on a public cloud because a hospital’s live health information management application processes and stores operational electronic protected health information (ePHI). This workload directly supports patient care and must preserve the confidentiality, integrity, and availability of sensitive health data. Although a public-cloud deployment can be made compliant, it requires a formal shared-responsibility design, an appropriate business associate agreement with the cloud provider, strong identity and access controls, encryption, audit logging, segmentation, and validated operational procedures. For a large, established application, the risk, governance effort, migration complexity, and potential impact of an outage or data exposure are greatest in production. Hospitals commonly evaluate less critical nonproduction environments first when considering cloud adoption, while retaining the live clinical workload in the environment that best satisfies their security, compliance, latency, and resilience requirements. The U.S. Department of Health and Human Services explains that covered entities must protect ePHI through administrative, physical, and technical safeguards under the HIPAA Security Rule. Palo Alto Networks also describes how cloud security requires clear responsibility for protecting workloads, configurations, identities, and data. See HHS HIPAA Security Rule and Palo Alto Networks: Cloud Shared Responsibility Model.

QUESTION NO: 5

Which three deployment modes of VM-Series firewalls are supported across NSX-T? (Choose three)

A. Partner Service

B. Boot Strap

C. Prism Central

D. Tier-1 insertion

E. Tier-0 insertion

ANSWER: A D E

Explanation:

VM-Series firewalls integrate with VMware NSX-T Data Center through the NSX-T Partner Service framework. This integration allows Panorama to register the VM-Series service with NSX-T Manager and automate deployment and lifecycle operations for the firewall service instances. The firewall can then be inserted at either supported logical-router tier, providing visibility and policy enforcement for traffic redirected through that router.

Tier-0 insertion places the VM-Series firewall on a tier-0 logical router, where it can secure traffic moving between the NSX-T logical network and physical or external networks. Tier-1 insertion places the firewall on a tier-1 logical router, enabling inspection and enforcement for traffic associated with tenant, application, or segment-level routing. In both insertion models, NSX-T redirects selected traffic to the firewall and the VM-Series appliance applies security policy received from Panorama. Palo Alto Networks documents Partner Service deployment and both tier-based insertion models as supported NSX-T deployment approaches. See the supported NSX-T deployments documentation and the VM-Series on NSX-T deployment guide.

QUESTION NO: 6

What is the major decision factor that customers use when selecting a managed container platform such as AS/EKS/GKE?

A. licensing costs

B. enhanced capabilities not available in vanilla K8s

C. no need to manage containers, just the application code.

D. reduced operational costs and management overhead

ANSWER: D

Explanation:

Reduced operational costs and management overhead is the primary decision factor for adopting a managed Kubernetes service such as Amazon EKS, Azure Kubernetes Service, or Google Kubernetes Engine. Kubernetes is powerful but operating its control plane, maintaining high availability, applying upgrades and security patches, managing etcd, and monitoring the underlying platform require substantial specialized effort. A managed service shifts much of this operational responsibility to the cloud provider, allowing the customer's platform and application teams to concentrate on workloads, deployment practices, and application delivery rather than on administering Kubernetes infrastructure.

This model can reduce the staffing time, operational risk, and ongoing maintenance associated with a self-managed Kubernetes environment. It also gives organizations access to provider-supported control-plane lifecycle management and integrations with cloud-native identity, networking, logging, and monitoring services. These operational advantages are especially important when teams need to deploy production container platforms quickly while maintaining reliability and consistent governance. AWS describes EKS as a managed Kubernetes service that removes the need to install, operate, and maintain a Kubernetes control plane; see AWS EKS documentation. Google similarly describes GKE as managing the Kubernetes control plane, as documented in the GKE overview.

QUESTION NO: 7

Which two statements describe Panorama template stacks? (Choose two.)

A. They combine settings from multiple templates for a managed firewall.

B. They support common baseline settings with site-specific configuration.

C. They centrally manage Security policy rulebases.

D. They define Log Collectors for centralized log storage.

ANSWER: A B

Explanation:

A template stack can **combine settings from multiple templates** into the configuration pushed to a managed firewall. This enables administrators to build modular configuration elements instead of creating a complete standalone template for every firewall or site.

Template stacks also support a **common baseline with location-specific or role-specific settings**. For example, an organization can use one template for shared DNS, NTP, and management settings and another template for site-specific interfaces, virtual routers, and zones. Device groups, not template stacks, are used to centrally manage policy rulebases. Panorama Collector Groups, not templates, manage centralized log collection. See the templates and template stacks documentation.

QUESTION NO: 8

How does the Palo Alto Networks NGFW integrate with Arista Networks Macro-

Segmentation Service?

A. Arista supports all hardware models of the Palo Alto Networks NGFW natively.

B. Arista allows standalone non-HA firewalls to be attached to a service leaf switch. You must configure an Elastic Load Balancer to obtain fault tolerance.

C. Arista CloudVision obtains relevant rules from Panorama through API and programs the Arista switches to steer intercepted east-west traffic to the Palo Alto Networks NGFW.

D. Arista owns the Security policy. It can extend the concept of fine-grained intra-hypervisor security for VMs by enabling dynamic insertion of services for virtualized devices such as firewalls

ANSWER: C

Explanation:

Arista CloudVision obtains relevant rules from Panorama through API and programs the Arista switches to steer intercepted east-west traffic to the Palo Alto Networks NGFW. This integration separates the security-policy and traffic-forwarding functions while coordinating them through APIs. Panorama remains the centralized management point for Palo Alto Networks security policy and the associated NGFW configuration. CloudVision uses the integration to obtain the information needed to identify traffic flows subject to inspection and to program the fabric for service insertion.

As a result, east-west flows can be transparently redirected through the NGFW, where App-ID, User-ID, Content-ID, and policy enforcement are applied, before traffic is forwarded toward its destination. This allows segmentation controls to be consistently applied in a high-speed data-center fabric without requiring endpoints or applications to be redesigned around firewall topology. The model combines CloudVision's fabric visibility and switch programming with the NGFW's Layer 7 security inspection and Panorama's centralized policy management. Palo Alto Networks describes its data-center architectures and NGFW-based segmentation capabilities at Palo Alto Networks Reference Architectures; Arista describes Macro Segmentation Service and its policy-driven service-insertion approach at Arista Macro Segmentation Service.