PassQueen
a.k.a dumpsqueen.com

VMware Cloud on AWS Master Specialist

VMware 5V0-11.21

Version Demo

Total Demo Questions: 8

Total Premium Questions: 85

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, VMware Cloud on AWS Architecture and Technologies	35
Topic 2, VMware Cloud on AWS Deployment and Configuration	26
Topic 3, VMware Cloud on AWS Management and Operations	22
Topic 4, VMware Cloud on AWS Troubleshooting and Optimization	2
Total	85

QUESTION NO: 1

Which two accounts are mandatory prerequisites for the successful deployment of a VMware Cloud on AWS solution? (Choose two.)

A. A VMware vCenter Server account

B. An AWS account

C. A VMware Cloud account

D. An Amazon Elastic Compute Cloud (EC2) account

E. A VMware Cloud on AWS account

ANSWER: B C

Explanation:

An AWS account and a VMware Cloud account are the required accounts for deploying a VMware Cloud on AWS software-defined data center. The VMware Cloud account, accessed through VMware Cloud services, provides the service entitlement and the console used to initiate and manage the SDDC deployment. Administrators use it to select the VMware Cloud on AWS service, choose the desired AWS region, configure the SDDC, and manage service-level settings.

The AWS account is required because the VMware Cloud on AWS SDDC is deployed on AWS infrastructure and must be associated with a customer-owned AWS account. During deployment, VMware Cloud on AWS establishes the necessary connection to AWS through an account link and requires permissions to create and manage the AWS-side networking and service resources that support the SDDC. These two accounts therefore establish both the VMware service-management identity and the AWS infrastructure ownership context needed for deployment.

VMware documents the AWS account linking and deployment prerequisites in the AWS VMware Cloud on AWS Getting Started Guide. VMware Cloud on AWS service documentation is also available from the VMware Cloud on AWS documentation portal.

QUESTION NO: 2

Refer to the exhibit.

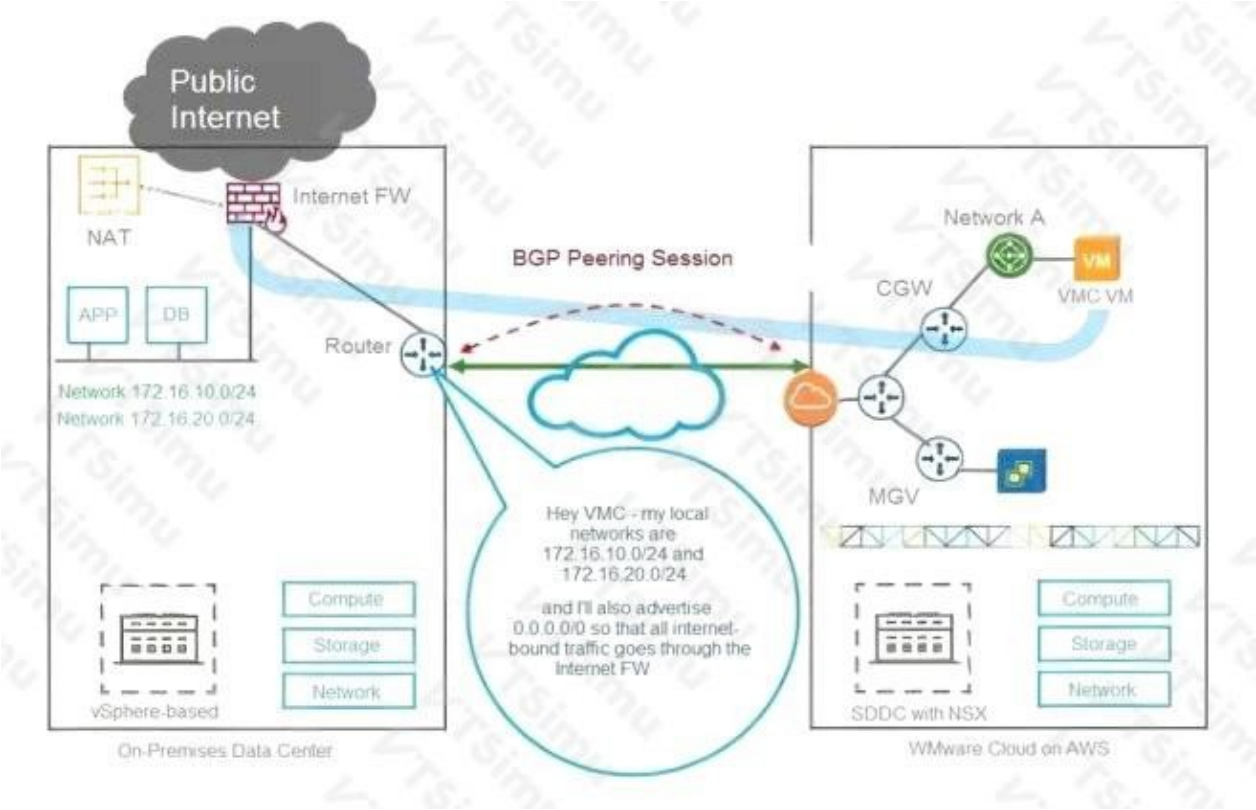

An administrator completes an assessment of its local data center for potential migration into VMware Cloud on AWS. After reviewing and analyzing the data and taking into account the company's business and IT priorities and budget constraints, connectivity between its on-premises and VMware Cloud on AWS environment will NOT require any high speed low latency connections. All new networks should be added to the software-defined data center (SDDC) routing table automatically when created. Which connection supports these requirements?

A. Layer 2 VPN (L2VPN)

B. AWS Direct Connect

C. Route-based VPN

D. Policy-based VPN

ANSWER: C

Explanation:

Route-based VPN is correct because it provides encrypted connectivity between the on-premises environment and VMware Cloud on AWS over the public Internet, making it appropriate when a dedicated high-bandwidth, low-latency connection is not required. A route-based VPN uses a virtual tunnel interface and can use dynamic routing with BGP. With BGP enabled, routes to on-premises and SDDC networks are exchanged dynamically rather than requiring administrators to manually maintain static encryption-domain or routing entries.

This dynamic route exchange is the key requirement in the scenario. When new routed networks are created in the SDDC, their prefixes can be advertised through BGP and learned by the on-premises routing infrastructure. Conversely, prefixes advertised from on premises can be learned by the SDDC gateway. This reduces operational overhead during migration and ongoing network expansion while retaining secure site-to-site connectivity. VMware Cloud on AWS networking documentation describes VPN connectivity and the use of route-based VPNs with dynamic routing. See the VMware Cloud on AWS documentation and the AWS Site-to-Site VPN documentation.

QUESTION NO: 3

An administrator is creating a runbook of the tasks that the Support team is responsible for within a new VMware Cloud on AWS software-defined data center (SDDC) configured with Hybrid Linked Mode.

Which two tasks must be performed by the Support team in both environments? (Choose two.)

A. Monitor workload virtual machines and infrastructure.

B. Patch guest operating systems.

C. Back up and restore workload virtual machines.

D. Manage the lifecycle of VMware vCenter Server.

E. Replace failed components.

ANSWER: B C

Explanation:

“Patch guest operating systems” and “Back up and restore workload virtual machines” remain customer/Support-team responsibilities in both the on-premises environment and the VMware Cloud on AWS SDDC. Hybrid Linked Mode provides unified visibility and inventory capabilities between the linked vCenter Server instances; it does not transfer administration of guest operating systems or workload-level data protection to VMware. The team responsible for the applications and virtual machines must therefore maintain guest OS security and update posture, including applying operating-system patches according to the organization’s change and vulnerability-management processes.

Likewise, workload VM backup and recovery must be designed, configured, operated, and validated by the customer. VMware Cloud on AWS protects and operates the managed cloud infrastructure, but customers remain responsible for selecting and managing a backup solution, defining retention and recovery requirements, and restoring their workload virtual machines when needed. VMware’s shared-responsibility guidance distinguishes VMware-managed SDDC infrastructure from customer-managed workloads and guest operating systems. See the VMware Cloud on AWS Operational Responsibilities documentation and the VMware Cloud on AWS shared responsibility guidance.

QUESTION NO: 4

Which two statements are true for the pre-defined resource pools in VMware Cloud on AWS? (Choose two.)

A. Users are allowed to monitor and modify the resource allocation settings in the Mgmt-RsourcePool.

B. The Mgmt-ResourcePool is able to utilize resources in subsequently created clusters if needed.

C. Users can modify the pre-configured vSphere DRS settings in their own Compute-ResourcePool.

D. Users can rename child resource pools to better match company policy.

E. All workload virtual machines are created in the top-level (root) Compute-ResourcePool by default.

ANSWER: D E

Explanation:

VMware Cloud on AWS provides a predefined resource-pool hierarchy that separates management workloads from customer workloads. Within the customer workload hierarchy, administrators can rename child resource pools to align their organization’s naming conventions and operational policies. This allows resource-pool names to clearly identify business units, applications, environments, or service tiers while retaining the supported VMware Cloud on AWS structure. Administrators can also create and organize workload resource pools beneath the Compute-ResourcePool as needed for allocation and organizational purposes.

All customer workload virtual machines are placed in the top-level, root Compute-ResourcePool by default when they are created. This default placement ensures that workload VMs run in the customer compute domain rather than the management resource pool, which is reserved for VMware-managed service components. If a different allocation or organizational model is required, an administrator can subsequently place workload VMs into an appropriate child resource pool under Compute-ResourcePool. VMware documents the predefined resource-pool layout and supported workload VM management behavior in its VMware Cloud on AWS guidance: VMware Cloud on AWS resource pool documentation and VMware Cloud on AWS documentation.

QUESTION NO: 5

An administrator would like their VMware Cloud on AWS software-defined data center (SDDC) cluster to scale down a host when CPU utilization drops below 60%. Which Elastic DRS policy should be selected?

A. Optimize for Lowest Cost

B. Optimize for Best Performance

C. Default Storage Scale-Out

D. Optimize for Rapid Scale-Out

ANSWER: A

Explanation:

Optimize for Lowest Cost is the appropriate Elastic DRS policy because it is designed to prioritize minimizing the number of hosts used by a VMware Cloud on AWS SDDC while retaining enough capacity for workload demand. Under this policy, Elastic DRS permits a more aggressive scale-in decision: when cluster utilization falls below the policy's scale-in threshold, it can recommend or perform removal of a host. The 60% CPU-utilization condition in the question aligns with the scale-in behavior associated with this cost-focused policy.

Elastic DRS continuously evaluates cluster CPU, memory, and storage capacity, along with policy objectives, to determine whether adding or removing hosts is appropriate. Selecting the lowest-cost policy is therefore suitable where the administrator's stated objective is to reduce host consumption once CPU demand declines to the specified level. This supports efficient SDDC resource consumption while Elastic DRS preserves the required capacity and operational constraints for the cluster. VMware documents Elastic DRS policy behavior and SDDC scaling operations in the VMware Cloud on AWS Operations Guide and provides service documentation through the VMware Cloud on AWS documentation portal.

QUESTION NO: 6

An organization has purchased both VMware Cloud on AWS and VMware vRealize Network Insight Cloud. Which additional integrated functionality are they able to utilize?

A. Creation of VMware HCX mobility groups from VMware vRealize Network Insight Cloud discovered applications

B. Automatic creation of network segment in VMware Cloud on AWS by VMware vRealize Network Insight Cloud

C. Collection of underlying AWS networking information sent to VMware vRealize Network Insight Cloud without additional configuration

D. Automatic software-defined data center (SDDC) grouping in VMware Cloud on AWS

ANSWER: A

Explanation:

Creation of VMware HCX mobility groups from VMware vRealize Network Insight Cloud discovered applications is correct. VMware vRealize Network Insight Cloud can discover applications and identify their communication dependencies, allowing an administrator to organize related workloads into migration waves. When VMware HCX is available in the environment, those discovered application waves can be exported to HCX as Mobility Groups through the HCX integration API. A Mobility Group provides a logical collection of virtual machines that should be migrated together, helping preserve application relationships and making migration planning more manageable.

This integration is especially useful for VMware Cloud on AWS migrations because it connects network- and application-dependency discovery with the HCX migration workflow. Rather than manually recreating workload groupings after analysis, the administrator can use the discovered application context to prepare migration groups for HCX operations. VMware documents the capability as exporting vRealize Network Insight application migration waves to HCX Mobility Groups. See the VMware HCX User Guide: vRealize Network Insight Integration and the VMware vRealize Network Insight Cloud documentation.

QUESTION NO: 7

What is a key functionality of the vRealize Automation Cloud Service Broker?

A. Provides a common catalog for easy consumption on VMware Cloud.

B. Manages blueprints as a code in a YAML format.

C. Automates the DevOps release lifecycle.

D. Creates and deploys virtual machines, applications, and services to multiple clouds.

ANSWER: A

Explanation:

Provides a common catalog for easy consumption on VMware Cloud. is correct because vRealize Automation Cloud Service Broker is the consumer-facing catalog component of VMware vRealize Automation Cloud (now VMware Aria Automation). It brings together approved cloud templates and extensibility actions into a unified service catalog, allowing users to discover and request resources through a consistent self-service experience. Administrators can organize catalog content, apply governance and approval policies, and make entitled items available to the appropriate projects and users. This model is particularly useful in hybrid and multi-cloud environments because consumers use one catalog rather than needing to understand the individual provisioning endpoints or underlying cloud implementations. Catalog items can include deployable templates as well as day-2 actions that enable users to perform permitted lifecycle operations after deployment. The key value is therefore simplified, governed consumption of services, with the catalog acting as the interface between cloud administrators and service consumers. VMware describes Service Broker as the service that provides and manages a catalog of cloud templates and other catalog items for consumption. See the VMware vRealize Automation Cloud documentation and the VMware Aria Automation product overview.

QUESTION NO: 8

Which two actions should an administrator complete before removing a host from a VMware Cloud on AWS cluster? (Choose two.)

A. Verify sufficient compute and memory capacity remains after evacuation.

B. Verify sufficient vSAN capacity and storage-policy compliance.

C. Delete all virtual machines on the selected host.

D. Disable vSphere HA for the cluster.

E. Detach the SDDC from its connected VPC.

ANSWER: A B

Explanation:

The administrator should verify that the remaining hosts have sufficient compute and memory capacity to accommodate the workload virtual machines currently running on the host selected for removal. During a supported scale-in operation, workload virtual machines are evacuated from the selected host and moved to hosts that remain in the cluster.

The administrator should also verify that vSAN has sufficient capacity and can maintain storage-policy compliance after the host is removed. Host removal requires vSAN data evacuation and can be blocked if the remaining cluster cannot satisfy the required availability policy or storage-capacity requirements. See the Scale a VMware Cloud on AWS SDDC cluster.