PassQueen
a.k.a dumpsqueen.com

VMware Cloud Professional

VMware 2V0-33.22

Version Demo

Total Demo Questions: 14

Total Premium Questions: 146

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, VMware Cloud Fundamentals	54
Topic 2, VMware Cloud on AWS	77
Topic 3, VMware Cloud Services	15
Total	146

QUESTION NO: 1

Which statement most accurately describes the service features of VMware Cloud on Dell EMC? (Select one option)

A. Dell technicians perform all software maintenance, as well as hardware fixes.

B. When an onsite response is required to fix a problem related to a host, a Dell technician must arrive on site within 24 hours.

C. An SDDC includes a minimum of one rack with three hosts. You can add hosts to the rack, up to the maximum supported by the rack.

D. VMwareSite Recovery is included as part of the initial service offering.

ANSWER: C

Explanation:

An SDDC includes a minimum of one rack with three hosts. You can add hosts to the rack, up to the maximum supported by the rack. This accurately describes the foundational deployment and scaling model for VMware Cloud on Dell EMC. The service delivered a VMware-managed software-defined data center on Dell EMC infrastructure installed in the customer's data center or edge location. A deployment begins with a single rack populated by the minimum supported host count of three hosts, providing the vSphere, vSAN, and NSX resources needed for the initial SDDC. Capacity can then be expanded non-disruptively by adding hosts to that same rack until the rack reaches its supported physical and service capacity. This model allows organizations to retain on-premises data locality while consuming the SDDC as a managed cloud service, with infrastructure capacity increased as workload requirements grow. VMware's launch material describes VMware Cloud on Dell EMC as an on-premises cloud service based on VMware Cloud Foundation software and Dell EMC infrastructure, while Dell's service announcement outlines its delivery in customer locations. See the VMware Cloud blog announcement and Dell Technologies service announcement.

QUESTION NO: 2

How much throughput does a Google Cloud VMware Engine private cloud network provide?

A. 25 Gbps

B. 40 Gbps

C. 100 Gbps

D. 10 Gbps

ANSWER: C

Explanation:

100 Gbps is correct because Google Cloud VMware Engine provides a dedicated VMware Engine network for a private cloud with 100-Gbps network throughput. This high-bandwidth network is part of the managed infrastructure that supports communication among VMware Engine resources and connectivity services associated with the private cloud. It is designed to provide the performance required for enterprise VMware workloads, including applications with substantial east-west traffic between virtual machines and infrastructure components. The throughput figure describes the capacity of the private cloud's VMware Engine network, rather than the effective throughput an individual workload will necessarily observe. Actual application performance can also depend on factors such as VM configuration, traffic patterns, connected services, and any applicable network design or policy controls. Google documents VMware Engine networking as a dedicated, high-throughput, low-latency networking capability and specifies the 100-Gbps private-cloud network capacity in its product documentation. See the Google Cloud VMware Engine networking overview and the VMware Engine overview for the service architecture and networking details.

QUESTION NO: 3

Which use cases apply to NSX logical routing? (Select two options)

A. You must provide external connectivity to VMs and containers.

B. Your organization must provide connectivity between VMs and containers that are connected to different segments.

C. You want to provide layer 2 connectivity between VMs and microservices.

D. You require intrinsic security for VMs connected to different segments.

ANSWER: A B

Explanation:

NSX logical routing is used to deliver Layer 3 connectivity within virtualized networks and from those networks to external destinations. “You must provide external connectivity to VMs and containers.” is correct because an NSX Tier-0 Gateway provides north-south routing between NSX workloads and physical networks. It can connect to upstream physical routers using dynamic routing protocols or static routes, enabling virtual machines and container-based workloads to reach networks beyond the NSX environment.

“Your organization must provide connectivity between VMs and containers that are connected to different segments.” is also correct because Tier-1 Gateways provide distributed routing between attached NSX segments. Workloads on separate segments use their configured default gateway addresses to communicate at Layer 3, and NSX performs this routing in the data path close to the workload. Tier-1 Gateways can also connect upward to a Tier-0 Gateway when traffic must reach external networks. Together, these gateway tiers support both east-west routing among isolated segments and north-south connectivity outside the NSX domain. See the NSX Tier-0 Gateway documentation and the NSX Tier-1 Gateway documentation.

QUESTION NO: 4

With which solution is the cloud administrator interfacing when defining storage policies in a VMware Cloud software-defined data center (SDDC)?

A. VMware Virtual Volumes (vVols)

B. VMware vSAN

C. iSCSI

D. VMware Virtual Machine File System (VMFS)

ANSWER: B

Explanation:

VMware vSAN is correct because VMware Cloud SDDCs use vSAN as their primary software-defined storage layer, and storage policies are defined through VMware’s Storage Policy-Based Management framework for vSAN-backed workloads. When an administrator creates or edits a VM storage policy, they select vSAN-specific rules that express the required level of availability, resilience, and performance for virtual-machine objects. For example, a policy can specify the number of failures a workload must tolerate, the storage architecture, and other vSAN capabilities. vSAN then applies that policy to the underlying objects and continuously monitors compliance, helping ensure that workloads receive the storage service level requested by the administrator.

In VMware Cloud on AWS, the vSAN datastore is presented as the SDDC’s shared workload storage. The administrator interfaces with vSAN capabilities when assigning storage policies to virtual machines or changing their storage requirements. This policy-driven approach is fundamental to vSAN because it connects a workload’s declared storage requirements to automated placement and protection of its data across the SDDC cluster. VMware’s documentation describes managing VM storage policies in a cloud SDDC and using policies for vSAN-backed virtual machines. See VMware Cloud on AWS documentation and VMware vSAN product information.

QUESTION NO: 5

A cloud administrator is tasked with improving the way that containers are scaled and managed in the environment. There is a currently no container orchestration solution implemented. Which solution can the administrator leverage to achieve this?

A. VMware NSX Container Plugin

B. Kubernetes

C. VMware vRealize Suite Lifecycle Manager

D. etcd

ANSWER: B

Explanation:

Kubernetes is the appropriate solution because it is a container orchestration platform designed to automate deployment, placement, scaling, and lifecycle management of containerized workloads. An administrator can define the desired number of application replicas, and Kubernetes continuously works to maintain that desired state across the available cluster nodes. This provides practical scaling through mechanisms such as Deployments, ReplicaSets, and autoscaling, while also handling workload scheduling, health monitoring, restart of failed containers, rolling updates, and service discovery.

For a cloud environment with no existing orchestration layer, adopting Kubernetes establishes the control plane and worker-node model needed to manage containers consistently at scale. It also provides declarative configuration: administrators describe the intended application state in manifests rather than manually starting, stopping, and relocating individual containers. VMware offerings such as Tanzu Kubernetes Grid build on Kubernetes to provide a supported Kubernetes runtime and lifecycle-management capabilities in VMware-based environments. Kubernetes documentation describes its role in automating deployment, scaling, and management of containerized applications, while VMware documents Tanzu Kubernetes Grid as a VMware-managed Kubernetes platform. See the Kubernetes overview and the VMware Tanzu Kubernetes Grid documentation.

QUESTION NO: 6

On VMware Cloud on AWS, which type of host do you use when you require high local storage requirements and additional cores for your workloads? (Select one option)

A. ve-standard-72

B. i3en.metal

C. i3.metal

D. AV36

ANSWER: B

Explanation:

i3en.metal is the appropriate VMware Cloud on AWS host type when workloads need both substantially higher local NVMe capacity and more CPU cores. This host family is designed for storage-intensive deployments, providing NVMe-based local storage that is consumed by VMware vSAN in a VMware Cloud on AWS SDDC. Compared with the earlier i3.metal-based host configuration, i3en.metal provides increased local storage capacity and additional physical CPU cores, making it well suited to data-heavy virtual machines, larger vSAN requirements, and applications that need greater compute density alongside low-latency local storage.

The underlying AWS i3en.metal instance provides 96 vCPUs, which correspond to 48 physical cores with hyperthreading, plus high-capacity NVMe SSD instance storage. In VMware Cloud on AWS, host selection must account for the storage and compute requirements of the cluster because the host-local disks contribute to the vSAN datastore. Therefore, selecting i3en.metal addresses the stated need for high local storage and additional cores in the same host platform. AWS documents the storage-optimized i3en instance family at AWS EC2 I3en Instances; VMware Cloud on AWS service information is available at VMware Cloud on AWS.

QUESTION NO: 7

A customer is looking to leverage a VMware Public Cloud solution to provide them with additional compute capacity as seasonal demand increases for their online business.

The current on-premises data center is configured as follows:

• VMware vSphere 7.0

• VMware vSphere Distributed Switch (vDS) 7.0

• Management and Server network - 172.18.0.0/16

• vMotion network - 192.168.120.0/24

• 250 application servers

Given the information in the scenario, which capability of VMware HCX will the customer not be able to utilize?

A. Cold migration

B. Layer 2 extension

C. Bulk migration

D. WAN optimization

ANSWER: D

Explanation:

WAN optimization is the HCX capability the customer cannot use with the stated on-premises environment. For the VMware HCX feature set and interoperability applicable to this scenario, the HCX WAN Optimization service does not support vSphere 7.0. WAN Optimization uses dedicated HCX WAN Optimization appliances to reduce the effects of constrained or high-latency WAN links through techniques such as data reduction, compression, and traffic optimization. Because the source environment is explicitly running vSphere 7.0, it does not meet the supported host-version requirement for deploying and using that service. This limitation applies independently of the number of application servers and the listed IP address ranges. Those details can affect migration planning, network-extension design, and the number of migration waves, but they do not make the unsupported WAN Optimization service available on vSphere 7.0. Before designing an HCX migration, administrators should validate the HCX release-specific compatibility requirements because supported versions and feature availability can change between HCX releases. VMware's HCX documentation is available through the VMware HCX documentation portal, and VMware provides the VMware Product Interoperability Matrix to verify supported product and version combinations.

QUESTION NO: 8

Which three components can be part of a virtual machine template? (Choose three.)

A. Installed applications, tools, and patches

B. vSphere tags

C. Custom attributes

D. Virtual Machine hardware configuration

E. Guest operating system

F. Virtual machine snapshots

ANSWER: A D E

Explanation:

A virtual machine template is a master image used to deploy consistent virtual machines. It preserves the virtual machine's hardware configuration, such as its virtual CPUs, memory, virtual disks, network adapters, firmware settings, and other configured virtual hardware. Consequently, deployments from the template begin with the standardized hardware profile defined by the template.

The template also includes the installed guest operating system. This enables administrators to create a prepared operating-system baseline rather than installing an operating system separately for every new virtual machine. The guest OS can be configured and generalized as appropriate before the template is used for deployment.

Installed applications, tools, and patches are part of the guest operating system's virtual disks and are therefore retained in the template. This makes templates particularly useful for delivering repeatable, policy-compliant workloads with required software, VMware Tools, and current updates already present. vSphere documentation describes templates as master copies that can include an installed guest operating system and applications, while the VM administration documentation details the virtual hardware configuration maintained for virtual machines and templates. See vSphere Virtual Machine Administration documentation and VMware virtual machine overview.

QUESTION NO: 9

A cloud administrator is tasked with deploying a new software-defined data center (SDDC) in VMware Cloud on AWS and has been able to log into the VMware Cloud console Successfully. However, they cannot access the VMware Cloud on AWS Services. Which two tasks need to be performed for the administrator to gain access? (Choose two.)

A. The cloud administrator will need to create a new subscription for the VMware Cloud on AWS service.

B. The cloud administrator will need to request access to the VMware Cloud on AWS service

C. The cloud administrator will need the globalcloudadmin role in the VMware Cloud on AWS service.

D. The cloud administrator will need the Administrator role in the VMware Cloud on AWS service.

E. The cloud administrator will need the cloudadmin role in the VMware Cloud on AWS service.

ANSWER: B D

Explanation:

Logging in to the VMware Cloud console authenticates the user, but it does not by itself grant entitlement to a VMware Cloud on AWS service or administrative permissions within that service. The cloud organization must first have access to VMware Cloud on AWS enabled through the service-access or subscription-request process. Therefore, "The cloud administrator will need to request access to the VMware Cloud on AWS service" is required when the organization has not yet been entitled to use the service.

After service access is available, the user also needs an appropriate service role. "The cloud administrator will need the Administrator role in the VMware Cloud on AWS service" grants the permissions needed to access and manage the VMware Cloud on AWS environment, including the ability to work with SDDCs and related service resources. VMware Cloud Console uses organization and service roles to control access; authentication alone is insufficient. The organization owner or another authorized administrator can assign the applicable VMware Cloud on AWS service role after the service has been activated or subscribed to. See VMware's VMware Cloud on AWS access documentation and the VMware Cloud services role documentation.

QUESTION NO: 10

VMware Engine cloud administrator is tasked with ensuring that a dedicated, secure, high-speed, and low-latency connection exists between an on-premises VMware Engine. Which two options are available for Google Cloud VMware Engine? (Choose two.)

A. Partner Interconnect

B. Global Reach

C. Dedicated Interconnect

D. ExpressRoute

E. Direct Connect

ANSWER: A C

Explanation:

Partner Interconnect and **Dedicated Interconnect** are the Google Cloud connectivity offerings available for establishing private connectivity from an on-premises environment to Google Cloud VMware Engine. Both services provide private network connectivity to Google's network rather than sending workload traffic over the public internet. This supports the security, performance, and predictable connectivity requirements commonly associated with hybrid VMware deployments.

Dedicated Interconnect is used when an organization needs direct physical connectivity between its on-premises network and Google's edge network. It is designed for high-capacity connectivity and is appropriate when the organization can establish connections at a supported Google colocation facility. Partner Interconnect provides comparable private connectivity through a supported service-provider partner. This lets an organization obtain connectivity where direct colocation access is impractical, while still extending its on-premises network privately into Google Cloud.

For Google Cloud VMware Engine, these interconnect products can be used with the required networking configuration to provide private, high-bandwidth paths between on-premises resources and the private cloud. Google documents both Dedicated Interconnect and Partner Interconnect as the two Cloud Interconnect connection types and describes their use for private hybrid connectivity. See the Cloud Interconnect overview and Google's VMware Engine networking overview.

QUESTION NO: 11

A cloud administrator Is managing a VMware Cloud on AWS environment consisting of a single cluster with six hosts. There have been no changes made to the Elastic DRS configuration.

In which two situations will Elastic DRS add another a host to the cluster? (Choose two.)

A. When availability zone failure occurs

B. When memory utilization reaches 90%

C. When network utilization reaches 90%

D. When CPU utilization reaches 90%

E. When storage utilization reaches 80%

ANSWER: A E

Explanation:

With the default Elastic DRS configuration, VMware Cloud on AWS automatically scales out the SDDC cluster when vSAN storage consumption reaches the configured storage utilization threshold. The default storage scale-out threshold is 80%, so reaching 80% storage utilization causes Elastic DRS to initiate the addition of a host. This proactive action provides additional vSAN capacity and maintains sufficient operational headroom as the environment grows.

Elastic DRS also adds a host when an AWS Availability Zone failure occurs. This is part of the service's availability and recovery behavior: VMware Cloud on AWS must restore the cluster's capacity and resilience after infrastructure loss in an Availability Zone. The host-addition operation is therefore not merely a response to ordinary workload demand; it supports recovery of the SDDC's intended resource and availability posture following the zone-level event.

These behaviors apply because the administrator has not changed Elastic DRS settings, so the documented default policy values and automated availability handling remain in effect. VMware documents Elastic DRS scale-out thresholds and Availability Zone failure behavior in the VMware Cloud on AWS Elastic DRS documentation. Additional operational guidance is available in the VMware Cloud on AWS documentation.

QUESTION NO: 12

A customer needs to set up a self-managed VDI solution that can be deployed to any VMware Cloud. Which two VMware solutions can meet this requirement? (Choose two.)

A. VMware Dynamic Environment Manager (DEM)

B. VMware ThinApp

C. VMware Workspace ONE Unified Endpoint Management (UEM)

D. VMware Horizon

E. VMware Workspace ONE Access

ANSWER: D E

Explanation:

VMware Horizon is the core VDI platform for delivering and managing virtual desktops and published applications. In a self-managed deployment, the customer operates the Horizon infrastructure, including its connection services, desktop pools, entitlements, and supporting VMware Cloud resources. Horizon is designed to support deployment across VMware-based cloud environments, allowing the organization to select the appropriate VMware Cloud platform while retaining administrative control of the VDI environment.

VMware Workspace ONE Access complements Horizon by providing identity and access capabilities for the self-managed end-user computing solution. It integrates with Horizon to provide a unified application catalog, identity federation, authentication policy, conditional access, and single sign-on to entitled virtual desktops and applications. Together, Horizon supplies the virtual desktop delivery layer and Workspace ONE Access supplies the secure identity and access layer needed for a complete, self-managed VDI service. The products are now maintained under the Omnissa brand, although they were historically sold and documented as VMware products. See the Omnissa Horizon product page and the Omnissa Workspace ONE Access product page.

QUESTION NO: 13

What is the purpose or the VMware cloud on AWS management gateway (MGW)?

A. A Tier-0 router that handles network traffic for workload virtual machines connected to routed computer network segments

B. A Tier-0 router that handles routing and firewalling for the VMware vCenter Server and other management appliances running in the software-defined datacenter (SDDC).

C. A Tier-1 router that handles network traffic for workload virtual machines connected to routes compute network segments

D. A Tier-1 router handles routing and firewalling for the VMware vCenter Server and Other management appliances running in the software-defined datacenter (SDDC).

ANSWER: D

Explanation:

The Management Gateway is a Tier-1 gateway in a VMware Cloud on AWS SDDC. Its purpose is to provide the routing and firewall enforcement point for the SDDC management network, which contains vCenter Server and other VMware management appliances. Management Gateway firewall rules determine which trusted source networks and addresses can reach management services. This separation lets administrators expose only the required management access while maintaining controlled connectivity to the SDDC's management components.

For a newly deployed SDDC, management access from the internet is not automatically available. Administrators configure Management Gateway firewall rules to permit the appropriate inbound access, such as access to vCenter Server, from explicitly approved sources. The gateway therefore plays an essential role in securely administering the cloud SDDC without mixing management-plane connectivity with the networking used by workload virtual machines. VMware's VMware Cloud on AWS networking and security documentation describes the Management Gateway as the Tier-1 router responsible for routing and firewalling of management appliances. See the VMware Cloud on AWS documentation and the Broadcom VMware Cloud on AWS technical documentation.

QUESTION NO: 14

Which two components are required in order to deploy a Tanzu Kubernetes Grid Cluster in VMware Could environment? (Choose two)

A. Tanzu CLI

B. Supervisor namespace

C. vSphere VM folder

D. vSphere resource pool

E. YAML manifest file

ANSWER: C D

Explanation:

A vSphere VM folder and a vSphere resource pool are required inventory targets for a Tanzu Kubernetes Grid deployment on vSphere-based cloud infrastructure. The VM folder provides the vCenter inventory location in which Tanzu Kubernetes Grid places the virtual machines that make up the management and workload clusters. Specifying this location keeps the deployed cluster VMs organized and allows vCenter inventory permissions and operational processes to be applied consistently.

The vSphere resource pool provides the compute placement target for those cluster virtual machines. Tanzu Kubernetes Grid uses the selected resource pool to place and run the VMs, ensuring that the cluster consumes compute resources from the intended vSphere hierarchy. These values are supplied as part of the vSphere deployment configuration, along with the applicable datacenter, datastore, network, and template details. In a VMware Cloud environment backed by vCenter, identifying both the folder and resource pool is therefore part of defining where the Tanzu Kubernetes Grid cluster is deployed and operated.

VMware's vSphere deployment guidance documents the required vCenter inventory configuration and placement settings for Tanzu Kubernetes Grid: VMware Tanzu Kubernetes Grid vSphere deployment guidance. See also the Tanzu Kubernetes Grid installation documentation.