

VMware Cloud Foundation Specialist (v2)

VMware 5V0-31.22

Version Demo

Total Demo Questions: 9

Total Premium Questions: 91

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, VMware Cloud Foundation Architecture and Technologies	12
Topic 2, VMware Cloud Foundation Planning and Design	28
Topic 3, VMware Cloud Foundation Installation, Configuration, and Setup	16
Topic 4, VMware Cloud Foundation Administration and Operations	33
Topic 5, VMware Cloud Foundation Troubleshooting and Repair	2
Total	91

QUESTION NO: 1

An administrator is planning a VMware Cloud Foundation VI workload domain that uses vSAN as principal storage.

Which two services should be available on every ESXi host participating in the vSAN cluster? (Choose two.)

- A. vSAN
- B. vMotion
- C. Fibre Channel boot from SAN
- D. NFS datastore connectivity
- E. SMB Direct

ANSWER: A B

Explanation:

vSAN must be configured on every host because vSAN uses local host storage devices to create the distributed vSAN datastore. Each participating host contributes storage resources and communicates with the other vSAN hosts across the vSAN network.

vMotion should also be configured on every host to provide workload mobility for operational activities, including maintenance and lifecycle workflows. VMware Cloud Foundation uses the workload-domain network configuration to apply the required vMotion service to cluster hosts. See the [VMware Cloud Foundation Administration Guide](#) and the [vSAN Administration documentation](#).

QUESTION NO: 2

An administrator is tasked with deploying a VMware Cloud Foundation environment that consists of three VI Workload Domains. Each VI Workload Domain is comprised of two clusters, with 18 hosts in each cluster.

Which option fulfills this requirement while minimizing the number of NSX-T Manager instances?

- A. Deploy one large-sized NSX-T Manager cluster for all VI Workload Domains.
- B. Deploy one medium-sized NSX-T Manager cluster for all VI Workload Domains.
- C. Deploy one medium-sized NSX-T Manager cluster per VI Workload Domain
- D. Deploy one large-sized NSX-T Manager cluster per VI Workload Domain.

ANSWER: B

Explanation:

Deploy one medium-sized NSX-T Manager cluster for all VI Workload Domains. The environment contains six clusters in total: two clusters per VI Workload Domain across three VI Workload Domains. With 18 hosts in each cluster, the NSX deployment must support 108 ESXi host transport nodes (six multiplied by 18). A medium NSX Manager appliance configuration supports up to 128 hypervisor transport nodes, so a single medium-sized NSX Manager cluster has sufficient capacity for all 108 hosts. This design also satisfies the requirement to minimize NSX Manager instances because the same NSX Manager cluster is shared rather than deploying a separate management cluster for each workload domain. NSX Manager sizing must account for the aggregate number of transport nodes managed by that NSX instance, not merely the number of workload domains or vSphere clusters. The medium deployment therefore provides the appropriate supported scale with the smallest number of NSX Manager clusters required for this topology. VMware documents NSX Manager appliance scale and transport-node limits in its [NSX Manager VM and Host Transport Node System Requirements](#). VMware Cloud Foundation planning guidance is available in the [VMware Cloud Foundation documentation](#).

QUESTION NO: 3

A systems administrator has recently added newly-commissioned hosts in the the VI workload domain, and IP addresses are automatically configured to their associated network pool. The administrator reviews which storage options require only vMotion and NFS networks in the network pool.

Which two storage options have this requirement? (Choose two.)

- A. vVols on iSCSI
- B. NFS
- C. vSAN and NFS
- D. vVols on NFS
- E. vSAN

ANSWER: B D

Explanation:

NFS and **vVols on NFS** are the storage choices that require the vMotion and NFS networks to be available in the VMware Cloud Foundation network pool for VI workload-domain hosts. The NFS network provides the IP connectivity from each ESXi host to the external NFS storage system. This same transport requirement applies when VMware Virtual Volumes are backed by an NFS storage container: host access to the protocol endpoint and the vVol datastore is provided through the NFS network.

The vMotion network is also required because VMware Cloud Foundation uses the network-pool definitions to apply the host networking needed for lifecycle and workload operations. When hosts are commissioned or expanded into a VI workload domain using either NFS-based option, SDDC Manager can allocate and configure the required IP addressing from the pool for both vMotion and NFS traffic. This allows workload mobility while retaining IP-based access to the NFS datastore infrastructure. VMware Cloud Foundation documents the supported workload-domain storage choices and their required network types in its deployment and planning guidance. See the [VMware Cloud Foundation documentation](#) and the [vSphere Virtual Volumes documentation](#).

QUESTION NO: 4

Which two configurations are part of the VMware Cloud Builder validation process? (Choose two.)

- A. License key Validates format, validity, and expiry for ESX, vSAN, vCenter Server, NSX, vRealize Suite, and Log Insight license keys
- B. Availability configuration: Validates the access to the configured backup locations
- C. Network configuration: Validates CIDR to IP address validity, IP addresses in use, gateways, invalid or missing VLANs. invalid or missing MTU, and network spec availability for all components
- D. Certificates: Validates certificates for ESX, vCenter Server, and NSX
- E. Passwords: Validates specified passwords Checks for minimum length, invalid characters, and format

ANSWER: A C E

Explanation:

VMware Cloud Builder performs pre-bring-up validation against the deployment parameters before it creates the initial VMware Cloud Foundation management domain. **License key Validates format, validity, and expiry for ESX, vSAN, vCenter Server, NSX, vRealize Suite, and Log Insight license keys** is included because Cloud Builder checks supplied license information to ensure that the required product licenses can be used for the deployment. **Network configuration: Validates CIDR to IP address validity, IP addresses in use, gateways, invalid or missing VLANs. invalid or missing MTU, and network spec availability for all components** is also a core validation area. Cloud Builder must verify addressing, routing, VLAN, MTU, and network-specification information because these settings enable management-domain components to communicate during bring-up. **Passwords: Validates specified passwords Checks for minimum length, invalid characters, and format** is likewise part of validation, ensuring the credentials in the deployment parameters meet

VMware Cloud Foundation password requirements before deployment begins. Although the prompt says to choose two, these three listed configuration checks are all part of the Cloud Builder validation process. VMware documents the Cloud Builder workflow and deployment validation as part of VMware Cloud Foundation bring-up in the [VMware Cloud Foundation documentation](#) and the [VMware Cloud Foundation documentation portal](#).

QUESTION NO: 5

A VMware Cloud Foundation administrator is required to enable Workload Management (vSphere with Tanzu) on an existing workload domain cluster, which is currently licensed with a vSphere Enterprise Plus license.

Which action, if any, is required to complete this task?

- A. Add a license for vSphere with Tanzu with sufficient CPU capacity to the SDDC Manager inventory, and then assign the license to the cluster in SDDC Manager
- B. No action is required since SDDC Manager licenses include an entitlement for vSphere with Tanzu.
- C. No action is required since the vSphere Enterprise Plus license supports vSphere with Tanzu.
- D. Add a license for vSphere with Tanzu with sufficient CPU capacity to both the SDDC Manager and vCenter Server, and then assign the license to the cluster in vCenter Server

ANSWER: A

Explanation:

Add a license for vSphere with Tanzu with sufficient CPU capacity to the SDDC Manager inventory, and then assign the license to the cluster in SDDC Manager. In VMware Cloud Foundation, SDDC Manager is the authoritative management plane for workload-domain licensing and for lifecycle-managed cluster configuration. Workload Management requires a vSphere with Tanzu entitlement; a standalone vSphere Enterprise Plus entitlement does not itself supply that entitlement. Before enabling the service, the administrator must import a vSphere with Tanzu license key that covers the physical CPU capacity of the target cluster into the SDDC Manager license inventory. The administrator then assigns that license to the workload-domain cluster through SDDC Manager. This assignment ensures the cluster is correctly licensed for Workload Management while preserving VMware Cloud Foundation's centralized inventory and configuration control. After the required Tanzu license is available and assigned, the administrator can proceed with the Workload Management enablement workflow for that cluster. VMware's VMware Cloud Foundation Kubernetes documentation describes adding the vSphere with Tanzu license in SDDC Manager and assigning it to the relevant cluster: [VMware Cloud Foundation documentation](#). VMware also documents the licensing prerequisites for vSphere with Tanzu in its [vSphere with Tanzu documentation](#).

QUESTION NO: 6

A service provider has a number of VMware Cloud Foundation workload domains and would like to sell Tanzu Namespaces as a managed service.

Which two functions will help the service provider with Tanzu resource management? (Choose two.)

- A. Object Limits
- B. Separate NSX-T instances
- C. Container Network Interfaces
- D. Resource Limits
- E. Resource Pools

ANSWER: A D

Explanation:

Object Limits and **Resource Limits** provide the core controls required to offer Tanzu Namespaces as a managed, multi-tenant service. Resource Limits let an administrator define the aggregate compute and storage capacity available to a namespace, including CPU, memory, and storage. This establishes an enforceable tenant boundary, enabling the provider to allocate capacity according to a service tier and prevent a namespace from consuming capacity reserved for other customers or workloads.

Object Limits complement those capacity quotas by controlling the number of Kubernetes objects that can be created in the namespace. A provider can therefore constrain the scale of tenant deployments, such as the permitted number of workload objects, while retaining predictable operational and infrastructure consumption. Together, these controls support chargeback-oriented service sizing, fair resource allocation, and protection of the underlying Supervisor and workload-domain resources. VMware's validated-design guidance identifies namespace resource limits and object limits as governance mechanisms for Tanzu namespace consumption. See the [VMware Validated Design documentation](#) and the [vSphere with Tanzu documentation](#).

QUESTION NO: 7

An administrator is preparing a vSAN cluster for a planned maintenance operation on one ESXi host. The administrator must ensure that all objects remain compliant with their assigned storage policies while the host is in maintenance mode.

Which vSAN data evacuation option should be selected?

- A. Ensure accessibility from other hosts
- B. Evacuate all data to other hosts
- C. No data migration
- D. Rebuild all objects after the host exits maintenance mode

ANSWER: B

Explanation:

Evacuate all data to other hosts is correct because this option migrates all accessible vSAN object components from the host entering maintenance mode to other hosts in the cluster. As a result, the objects remain compliant with their assigned storage policies while the host is unavailable.

The Ensure accessibility from other hosts option maintains object accessibility but can leave objects with reduced redundancy. No data migration leaves components on the host and is intended only for short maintenance events where the host is expected to return quickly. See the [vSAN Administration documentation](#).

QUESTION NO: 8

An architect is presented with a Bill of Material for different vSAN ready nodes:

- 8 identical nodes with 1 vSAN disk group
- 8 identical nodes with 2 vSAN disk groups

What is the ideal design of VMware Cloud Foundation domains that the architect should recommend?

- A. 1 Management Domain of 4 nodes, 1 Workload Domain of 4 nodes, and 1 Workload Domain of 8 nodes
- B. 1 Domain of 16 nodes deployed on a Consolidated Architecture model
- C. 1 Management Domain of 6 Nodes, 1 Workload Domain of 4 nodes, and 1 Workload Domain of 6 nodes
- D. 1 Management Domain of 6 nodes and 1 Workload Domain of 10 nodes

ANSWER: A

Explanation:

1 Management Domain of 4 nodes, 1 Workload Domain of 4 nodes, and 1 Workload Domain of 8 nodes is the appropriate design because it respects the two distinct ReadyNode hardware configurations while creating a supported, practical domain layout. The eight nodes with one vSAN disk group can be divided into a four-node management domain and a four-node workload domain. A four-node management domain provides the standard resilient footprint for the management components and vSAN storage, while the remaining four matching nodes form a separate workload domain without introducing dissimilar disk-group configurations into the same vSAN cluster.

The eight identical nodes with two vSAN disk groups should remain together as an eight-node workload domain. Keeping these nodes together preserves hardware consistency within that vSAN cluster and makes their greater storage-device parallelism available to workloads that require additional performance or capacity. This also produces clear operational boundaries: management services run in their dedicated management domain, and workload requirements can be placed in the workload domain whose ReadyNode profile best fits them. VMware Cloud Foundation uses workload domains as independent deployment and lifecycle-management units, and vSAN ReadyNodes are validated configurations intended to be deployed consistently in a vSAN cluster. See the [VMware Cloud Foundation documentation](#) and the [VMware Compatibility Guide for vSAN ReadyNodes](#).

QUESTION NO: 9

An administrator is tasked with providing additional North-South throughput to the workloads hosted on overlay-backed networks in a VI Workload Domain stretched cluster. A two-node NSX Edge cluster was previously deployed through SDDC Manager before the cluster was stretched.

Which option is valid to add two nodes to the existing edge cluster while maintaining password rotation capability?

- A. Expand the existing NSX Edge cluster using NSX Manager, and import the additional nodes in SDDC Manager.
- B. Create a new NSX Edge cluster using SDDC Manager on the same stretched Workload Domain cluster.
- C. Expand the existing NSX Edge cluster using SDDC Manager, and place the additional NSX Edge nodes on the same stretched Workload Domain cluster.
- D. Expand the existing NSX Edge cluster using SDDC Manager, and place the additional NSX Edge nodes on a different stretched Workload Domain cluster.

ANSWER: C

Explanation:

Expand the existing NSX Edge cluster using SDDC Manager, and place the additional NSX Edge nodes on the same stretched Workload Domain cluster. is valid because SDDC Manager remains the lifecycle-management authority for an NSX Edge cluster deployed as part of a VMware Cloud Foundation workload domain. Expanding the Edge cluster through SDDC Manager adds the required Edge appliances while preserving the inventory, configuration, and credential-management relationship between SDDC Manager and NSX Manager. This is essential for continued use of the VMware Cloud Foundation password-management workflow, including supported password rotations for the NSX Edge node components.

Adding two Edge nodes increases the Edge cluster's available data-plane capacity, allowing additional North-South traffic processing for overlay-backed workloads. Because the workload domain has been stretched, the additional Edge nodes should be deployed in that same stretched workload domain so that their placement aligns with the domain's availability and site-topology design. SDDC Manager performs the supported deployment orchestration, including the required network and host selections, and records the resulting configuration for future lifecycle operations. VMware Cloud Foundation administration documentation is available at [VMware Cloud Foundation Documentation](#); NSX Edge architecture and deployment concepts are documented at [VMware NSX Documentation](#).