

## VMware New 2023 5V0-31.22 Sample Questions Reliable 5V0-31.22 Test Engine [Q16-Q34]



## VMware New 2023 5V0-31.22 Sample Questions Reliable 5V0-31.22 Test Engine [Q16-Q34]

VMware New 2023 5V0-31.22 Sample Questions Reliable 5V0-31.22 Test Engine  
Feel VMware 5V0-31.22 Dumps PDF Will likely be The best Option

The VMware 5V0-31.22 certification exam is designed to validate the skills and knowledge of professionals in VMware Cloud Foundation technology. This certification is intended for individuals who have experience in deploying and managing VMware Cloud Foundation solutions. The exam covers various topics such as installation, configuration, management, troubleshooting, and optimization of VMware Cloud Foundation components. Passing this exam indicates that the candidate has a solid understanding of VMware Cloud Foundation and is equipped with the necessary skills to deploy and manage these solutions.

### QUESTION 16

The vSAN Witness appliance for a VMware Cloud Foundation stretched vSAN cluster stopped working. The administrator needs to roll out a new appliance to replace the old one.

Which tool should the administrator use to perform this task?

- \* vSphere Client
- \* SDDC Manager
- \* vSphere Update Manager
- \* vSAN PowerCLI

Explanation

This is because according to VMware documentation, this is the tool that an administrator should use to perform this task of rolling out a new vSAN Witness appliance for a VMware Cloud Foundation stretched vSAN cluster. SDDC Manager provides a user interface and API for deploying and managing VCF components, including vSAN Witness appliances.

### QUESTION 17

Which two options are only available when using vSphere Lifecycle Manager Images? (Choose two.)

- \* Upgrade VM Hardware Compatibility versions.
- \* Update the firmware of all ESXi hosts in a cluster.
- \* Install and update third-party software on all ESXi hosts in a cluster.
- \* Check the hosts and clusters against the vSAN Hardware Compatibility List.
- \* Upgrade and patch ESXi hosts.

Explanation

This is because vSphere Lifecycle Manager images can include firmware updates and third-party software components that can be applied to all hosts in a cluster<sup>12</sup>. These options are only available when using vSphere Lifecycle Manager images, not when using vSphere Lifecycle Manager baselines<sup>2</sup>.

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere-lifecycle-manager.doc/GUID-9A11223>

### QUESTION 18

What is a characteristic about the Credentials Worksheet in the Deployment Parameter Workbook?

- \* Passwords can be common only for appliance users.
- \* Passwords can be different per user.
- \* Passwords must be different per user.
- \* Passwords must be common across all users.

Explanation

According to VMware Cloud Foundation Planning and Preparation Workbook, when filling out the Credentials Worksheet in the Deployment Parameter Workbook, you must provide different passwords for each user account that will be created during deployment. This ensures security and compliance for your environment.

### QUESTION 19

Which license is required to enable Workload Management on VMware Cloud Foundation?

- \* VMware vSphere Evaluation
- \* VMware vSphere Standard
- \* VMware vSphere Enterprise Plus
- \* VMware Tanzu Basic

Explanation

A Tanzu Basic license is required to enable Workload Management on VMware Cloud Foundation. Once enabled, the Supervisor Cluster must be assigned a Tanzu license before the 60-day evaluation period expires.

This license can be added to the license inventory of vSphere if a valid Tanzu Edition license is available.

## QUESTION 20

An administrator has registered an external identity source in a consolidated architecture and would like to make sure that any subsequent workload domains can be accessed using the same identity sources.

How can this goal be achieved with VMware Cloud Foundation?

- \* By configuring IWA as an identity source
- \* By configuring LDAPS as an identity source
- \* By keeping the pre-configured defaults
- \* By replicating vSphere SSO configuration

Explanation

vSphere Single Sign-On (SSO) provides secure authentication and authorization services for VMware Cloud Foundation components, including vCenter Server and Platform Services Controller (PSC). In a consolidated architecture deployment of VMware Cloud Foundation, the vSphere SSO configuration is shared across all the workload domains.

To ensure that subsequent workload domains can use the same identity sources as an external identity source registered in a consolidated architecture, the administrator needs to replicate the vSphere SSO configuration.

This can be achieved by configuring the same identity sources for vSphere SSO across all the workload domains.

Configuring IWA (Integrated Windows Authentication) or LDAPS (Lightweight Directory Access Protocol over SSL) as an identity source is a part of configuring the vSphere SSO configuration for identity sources.

Keeping the pre-configured defaults does not guarantee that the subsequent workload domains will use the same identity sources as the external identity source registered in a consolidated architecture.

References:

- \* VMware Cloud Foundation Operations and Administration

Guide:<https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html>

- \* VMware vSphere Security

Guide:<https://docs.vmware.com/en/VMware-vSphere/7.0/vsphere-security-guide.pdf>

\* To ensure that any subsequent workload domains can be accessed using the same identity sources, it is necessary to replicate the vSphere SSO configuration across all the workload domains in a consolidated architecture deployment. This can be achieved by replicating the vSphere SSO configuration between the primary and additional SDDC Manager instances. This ensures that all the workload domains registered with the SDDC Manager will be able to consume resources and services from the same identity sources without any additional configuration in each individual workload domain.

## QUESTION 21

A systems administrator is tasked to deploy a management domain during VMware Cloud Foundation Bring-Up process. What are the minimum hardware requirements for the management cluster?

- \* 2 vSAN Ready Nodes, 192 GB RAM per server, and 4 10GbE NICs
- \* 8 vSAN Ready Nodes, 256 GB RAM per server, and 2 10GbE NICs
- \* 4 vSAN Ready Nodes, 192 GB RAM per server, and 2 10GbE NICs
- \* 6 vSAN Ready Nodes, 256 GB RAM per server, and 4 10GbE NICs

Explanation

The minimum hardware requirements for the management cluster during VMware Cloud Foundation Bring-Up process are:

- \* 4 vSAN Ready Nodes
- \* 192 GB RAM per server
- \* 2 10GbE NICs

## QUESTION 22

A VMware Cloud Foundation administrator created a Tanzu Namespace in one of the workload domains.

Which two functions related to permissions can be performed on the newly created Namespace? (Choose two)

- \* Add permissions only from the vSphere.local domain.
- \* Permissions can be set to either view or edit.
- \* Add permissions to users from vCenter Single Sign-On identity sources.
- \* Add a custom role to create more granular permissions.
- \* Add permissions to local vSphere with Tanzu users only.

Explanation

A quote from reference [1] states that, "To add permissions to users or groups from vCenter Single Sign-On identity sources, the Tanzu Kubernetes cluster administrator can use either the vSphere Client or kubectl." Another quote from reference [1] states that, "By default, a Tanzu Kubernetes cluster includes a set of predefined roles that provides granular permission control for Kubernetes objects. The predefined roles enable cluster groups to be created with specific permissions across the Kubernetes namespace hierarchy.

Administrators can also create custom roles to provide more granular permission control that is specific to their organization's requirements. References: [1] Tanzu Kubernetes Cluster or Supervisor Cluster[1]: Which do I choose?

<https://blogs.vmware.com/virtualblocks/2022/06/23/tanzu-kubernetes-cluster-or-supervisor-cluster-which-do-i-c>

## QUESTION 23

Which order of steps should an administrator use to replace a failed host in a stretched cluster?

- \* Decommission the failed host.
2. Remove the host using cluster APIs.
  3. Add the newly commissioned host to the cluster using cluster APIs. 4 Commission the new host with the correct network.
- \* 1 Remove the host using cluster APIs.
2. Decommission the failed host.

3. Commission the new host with the correct network.
4. Add the newly commissioned host to the cluster using cluster APIs.
  - \* Remove the host using cluster APIs
2. Decommission the failed host.
3. Add the newly commissioned host to the cluster using cluster APIs.
4. Commission the new host with the correct network
  - \* 1 Decommission the failed host
2. Remove the host using cluster APIs.
3. Commission the new host with the correct network.
4. Add the newly commissioned host to the cluster using cluster APIs.

Explanation

This is because according to VMware documentation<sup>1</sup>, these are the steps to replace a failed host in a stretched cluster:

- \* Run the compact cluster API to remove any stale data from vSAN.
- \* Decommission the host to be removed using SDDC Manager UI or API.
- \* Commission the replacement host to the same network pool as the removed host using SDDC Manager UI or API.
- \* Add the newly commissioned host to the cluster using SDDC Manager UI or API.
- \* Explanation: According to the VMware documentation, when replacing a failed host in a stretched cluster, the first step is to decommission the failed host. This should be followed by removing the host using cluster APIs, commissioning the new host with the correct network, and then adding the newly commissioned host to the cluster using cluster APIs.

## QUESTION 24

Which two options can be used to create a new VMware Cloud Foundation VI workload domain? (Choose two.)

- \* SDDC Manager UI
- \* PowerCLI
- \* Cloud Builder UI
- \* vCenter UI
- \* REST API

Explanation

The SDDC Manager UI provides a single point of control for managing and monitoring your VMware Cloud Foundation instance and for provisioning workload domains. You use the navigation bar to move between the main areas of the user interface <sup>1</sup>. The SDDC Manager UI provides an integrated view of the physical and virtual infrastructure and centralized access to manage the physical and logical resources <sup>2</sup>.

The REST API can also be used to create a new VI workload domain using VMware Cloud Foundation. The VMware Cloud Foundation API Reference Guide provides information on available operations <sup>3</sup>.

## QUESTION 25

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.)

- \* Object Limits
- \* Separate NSX-T instances
- \* Container Network Interfaces
- \* Resource Limits
- \* Resource Pools

Explanation

This is because according to VMware documentation, these are some of the functions that will help the service provider with Tanzu resource management:

- \* Object Limits: You can specify limits on objects such as pods, services, persistent volume claims, etc.

for each namespace.

- \* Resource Limits: You can specify limits on resources such as CPU and memory for each namespace.

## QUESTION 26

During a VCF design workshop, the architect gathered the following customer requirements:

- \* There must be two environments: PROD and DEV.
- \* PROD and DEV should be administratively separated.
- \* PROD will use two different hardware server types, and DEV will only use one hardware server type.
- \* The VCF infrastructure design should be flexible and scalable as much as possible How many NSX local managers in total will be provisioned after deploying the full VCF infrastructure?
  - \* 6
  - \* 3
  - \* 12
  - \* 9

Explanation

According to the VMware documentation, each NSX-T Local Manager is associated with a vCenter Server, and each NSX-T Local Manager can manage up to three vCenters. In a VCF deployment with two environments (PROD and DEV) and two different hardware server types in PROD, there would be a total of three vCenter Servers. Therefore, a total of three NSX-T Local Managers would be provisioned to manage the three vCenter Servers.

## QUESTION 27

Which two roles are played by a Spherelet in a Tanzu-enabled VCF workload domain? (Choose two.)

- \* It runs as a VIB on all Supervisor Cluster ESXi hosts configured with the vSphere Networking Stack.

- \* It enables an ESXi hypervisor to act as a Kubernetes master node.
- \* It enables an ESXi hypervisor to act as a Kubernetes worker node.
- \* It starts and monitors vSphere pods running on the workload domain cluster
- \* It communicates with the vSphere with Tanzu embedded Harbor registry.

Explanation

According to vSphere with Tanzu Architecture<sup>1</sup>, a Spherelet is a component that runs as a VIB on all Supervisor Cluster ESXi hosts configured with the vSphere Networking Stack. It enables an ESXi hypervisor to act as a Kubernetes master node or a Kubernetes worker node, depending on the role assigned by the Supervisor Cluster control plane.

## QUESTION 28

A VCF architect collected the following requirements when designing the expansion of a new VI Workload Domain with twenty four vSAN Ready nodes, each with a dual-port 25Gbps network interface card:

- \* Provide scalable high-performance networking with layer-3 termination at top-of-rack
- \* Protect workloads from switch/NIC/rack failure
- \* Provide isolation for DMZ workloads
- \* Provide at-least 25Gbps dedicated bandwidth to backup traffic
- \* Easily accept workloads on traditional VLAN-backed networks
- \* Fully-supported by VMware

Which three design considerations meet all of these requirements? (Choose three.)

- \* Two-node Edge Cluster with ECMP
- \* Spine and Leaf network topology with layer-3 at Spine
- \* Stretched Clustering
- \* Spine and Leaf network topology with layer-3 at top of rack
- \* Two-node Edge Cluster with BFD
- \* Core Aggregation network topology

Explanation

Option B: Spine and Leaf network topology with layer-3 at Spine &#8211; A spine and leaf network topology is designed for high scalability and performance, and layer-3 at the spine ensures that there is no single point of failure for the layer-3 termination. This meets several of the requirements, including scalable high-performance networking with layer-3 termination at top-of-rack, protecting workloads from switch/NIC/rack failure, and providing isolation for DMZ workloads.

Option D: Spine and Leaf network topology with layer-3 at top of rack &#8211; Similar to Option B, this topology also provides high scalability and performance, and layer-3 at the top of rack meets the requirement for layer-3 termination at top-of-rack.

Option F: Core Aggregation network topology &#8211; This topology provides a highly available, redundant core switch for aggregation and routing, which meets the requirement for protecting workloads from switch/NIC/rack failure.

Based on the given choices, the correct answers would be B, D, and F.

Sources: [1] Designing VMware Infrastructure Topology and Architecture; Authors: Russel Nolan, Eiad Al-Aqqad [2] Network

## Topology Considerations for VMware

vSAN; <https://docs.vmware.com/en/VMware-vSAN/7.0/com.vmware.vsan.networking.doc/GUID-1A901C10-48> Spine-Leaf Architecture:

Introduction; <https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c>

### QUESTION 29

Which three components are required to deploy a stretched cluster in a VMware Cloud Foundation environment? (Choose three.)

- \* vSAN, host overlay and vMotion network stretched across both sites
- \* DHCP on the NSX Edge overlay network
- \* DHCP on the host overlay network
- \* One witness host per site
- \* One witness host per vSAN stretched cluster
- \* vSAN: host overlay and vMotion network per data site

#### Explanation

This is because when deploying a stretched cluster in a VMware Cloud Foundation environment, you must ensure that:

- \* The vSAN network, the host overlay network and the vMotion network are stretched across both availability zones<sup>12</sup>. These networks enable data replication, communication and migration between hosts in different sites.
- \* The host overlay network has DHCP enabled to provide IP addresses to hosts<sup>1</sup>. This simplifies the configuration and management of hosts in different sites.
- \* Each vSAN stretched cluster has one witness host deployed in a separate location from both availability zones<sup>12</sup>. The witness host acts as a tie-breaker in case of a site failure or split-brain scenario.

<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-admin/GUID-7B4CC729-20BD-4CC9-B855-B>

### QUESTION 30

An administrator wants to delete a VMware Cloud Foundation Workload Domain and re-use the attached ESXi hosts by returning them to the list of unassigned hosts in the SDDC Manager inventory.

Which action needs to be taken to complete this task?

- \* ESXi hosts need to be re-imaged and updated.
- \* ESXi hosts need to be decommissioned and re-imaged.
- \* ESXi hosts need to be re-imaged and rejoined.
- \* ESXi hosts need to be decommissioned and updated

#### Explanation

This is because according to VMware documentation, this is the procedure for deleting a VMware Cloud Foundation Workload Domain and re-using its ESXi hosts:

- \* Decommission all ESXi hosts in a cluster
- \* Delete all clusters in a workload domain
- \* Delete workload domain



- \* Re-image ESXi hosts using SDDC Manager

### QUESTION 31

A customer purchased six new HPE ProLiant DL380 Gen10 hosts and is interested in deploying a VCF infrastructure that will coexist with the current VMware on AWS cloud DR solution in a hybrid model.

The architect suggests the VCF architecture that will help the customer run the workloads while offering workload isolation.

Which VCF architecture was suggested for this customer?

- \* node management domain and 3-node VI workload domain in a VCF standard model, while leveraging resource pools
- \* 3-node management domain and 3-node VI workload domain in a VCF standard model
- \* node management/workload domain in a VCF consolidated model, while leveraging resource pools
- \* node management/workload domain in a VCF consolidated model

Explanation

This is because according to VMware documentation, VCF supports two architecture models; standard and consolidated. The standard architecture model separates management workloads and user workloads into different domains, while the consolidated architecture model combines them into one domain. For a hybrid model that coexists with VMware on AWS cloud DR solution, the standard architecture model is recommended as it provides workload isolation and mobility across VCF instances<sup>6</sup>. The minimum requirement for creating a management domain or a VI workload domain is three hosts<sup>4,5</sup>.

### QUESTION 32

A systems administrator is implementing stretched clusters in an environment with multiple Availability Zones (AZs). Which statement accurately describes this design?

- \* If VLAN is stretched between AZ1 and AZ2, the Layer 3 network must also be stretched between the two AZs.
- \* Layer 3 networks must be stretched between the AZs by the physical infrastructure
- \* The Layer 3 gateway for the workload domain and Edge overlay networks must be highly available across the AZs.
- \* For VLANs that are stretched between AZs, configure load balancing in the Layer 3 gateway between AZs

Explanation

This is because according to VMware documentation, this is one of the design considerations for implementing stretched clusters in an environment with multiple Availability Zones (AZs). Load balancing in the Layer 3 gateway between AZs can improve network performance and availability by distributing traffic across multiple paths.

### QUESTION 33

Which two options can be used to create a new VMware Cloud Foundation VI workload domain? (Choose two.)

- \* PowerCLI
- \* REST API
- \* Cloud Builder UI
- \* vCenter UI
- \* SDDC Manager UI

Explanation

The SDDC Manager UI provides a single point of control for managing and monitoring your VMware Cloud Foundation instance and for provisioning workload domains. You use the navigation bar to move between the main areas of the user interface <sup>1</sup>. The SDDC Manager UI provides an integrated view of the physical and virtual infrastructure and centralized access to manage the

physical and logical resources 2.

The REST API can also be used to create a new VI workload domain using VMware Cloud Foundation. The VMware Cloud Foundation API Reference Guide provides information on available operations 3.

#### QUESTION 34

A developer is deploying pods with Persistent Volumes (PV) on vSphere with Tanzu. Which component determines the datastore that the PV will be placed on?

- \* CNS-CSI
- \* Hostd
- \* Spherelet
- \* SPBM

Explanation

This is because according to VMware documentation<sup>34</sup>, vSphere with Tanzu uses storage policies to integrate with shared datastores available in your environment, including VMFS, NFS, vSAN, or vVols datastores. The storage policies represent datastores and manage the storage placement of such objects as persistent volumes (PVs). Storage Policy Based Management (SPBM) is a framework that provides a single unified control plane across different types of datastores and enables administrators to define policies based on storage capabilities and requirements<sup>5</sup>.

To prepare for the VMware 5V0-31.22 exam, VMware offers several resources, including official courses, study guides, and practice exams. Candidates can also gain hands-on experience with VMware Cloud Foundation by deploying and managing the solution in a lab environment. By earning the VMware Cloud Foundation Specialist (v2) certification, IT professionals can demonstrate their expertise in this advanced technology and enhance their career opportunities in the IT industry.

The VMware Cloud Foundation Specialist (v2) exam is a comprehensive assessment of a candidate's knowledge and expertise in VMware Cloud Foundation. The exam is divided into three sections: Architecture, Deployment, and Troubleshooting. The Architecture section covers the design and architecture of VMware Cloud Foundation, including the components and their functionality. The Deployment section focuses on the installation, configuration, and upgrade of VMware Cloud Foundation components. Finally, the Troubleshooting section tests the candidate's ability to identify and resolve issues with VMware Cloud Foundation solutions.

**Use Valid New 5V0-31.22 Test Notes & 5V0-31.22 Valid Exam Guide:**

<https://www.examslabs.com/VMware/VMware-Certified-Specialist/best-5V0-31.22-exam-dumps.html>