

Vmware.Premium.5V0-31.22.70q - DEMO

Number: 5V0-31.22  
Passing Score: 800  
Time Limit: 120 min  
File Version: 1.7



**Exam Code: 5V0-31.22**

**Exam Name:** VMware Cloud Foundation Specialist (v2)

**Website:** [www.VCEplus.io](http://www.VCEplus.io) - [www.VCEup.com](http://www.VCEup.com)

VCEup

## Exam A

### QUESTION 1

An administrator needs additional capacity on a vSAN cluster. Each host currently has only one disk group. Which two approaches can be used to expand storage capacity in this situation? (Choose two.)

- A. Increase the number of cache disks in the existing disk group.
- B. Add an additional disk group.
- C. Disable compression.
- D. Increase the number of capacity disks in the existing disk group
- E. Disable deduplication.

**Correct Answer: B, D**

**Section:**

**Explanation:**

To expand storage capacity in a vSAN cluster with one disk group, you can either add more drives to hosts in the cluster, which is commonly referred to as scaling up, or add capacity drives to existing disk groups

Option B: Add an additional disk group - According to search result [1], adding additional drives to a host will increase both capacity and performance [1], and each disk group contains one flash cache device and one or multiple capacity devices for persistent storage [2]. Therefore, adding an additional disk group to each host would increase the storage capacity of the vSAN cluster.

Option D: Increase the number of capacity disks in the existing disk group - Search result [1] explains that vSAN clusters require capacity and cache devices to function, and each disk group can contain multiple capacity devices for persistent storage [2]. Thus, an additional way to expand storage capacity in the vSAN cluster would be to increase the number of capacity disks in the existing disk group.

Reference: 1: VMware vSAN documentation 2: VMware vSAN documentation

A disk group is a collection of one or more flash-based cache devices and one or more capacity devices that provide storage capacity for a vSAN cluster. A vSAN cluster can have multiple disk groups, and each disk group can have a different configuration.

To expand storage capacity in a vSAN cluster where each host currently has only one disk group, the administrator can add an additional disk group or increase the number of capacity disks in the existing disk group.

Adding an additional disk group involves adding more disks to the host and creating a new disk group. This approach can provide additional capacity and performance benefits, as the new disk group can be configured with different settings to optimize performance and capacity.

Increasing the number of capacity disks in the existing disk group involves adding more capacity devices to the existing disk group. This approach can provide additional capacity, but may not necessarily provide performance benefits as the existing disk group may already be fully utilized.

Reference:

VMware vSAN 7.0 Design and Sizing Guide: <https://storagehub.vmware.com/t/vmwarevsan/vmware-vsan-7-0-design-and-sizing-guide-2/VMware> vSAN Documentation:

<https://docs.vmware.com/en/VMwarevSphere/7.0/com.vmware.vsphere.vsan-planning.doc/GUID-9B7C9685-64C5-49C2-8E3CCC2E47AFBC6F.html>

### QUESTION 2

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.)

- A. Two-node Edge Cluster with ECMP
- B. Spine and Leaf network topology with layer-3 at Spine
- C. Stretched Clustering
- D. Spine and Leaf network topology with layer-3 at top of rack
- E. Two-node Edge Cluster with BFD
- F. Core Aggregation network topology

**Correct Answer: B, D, F**

**Section:**

**Explanation:**

Option B: Spine and Leaf network topology with layer-3 at Spine - 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 - 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 - 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-1A901C10-4894-4E9B-8A36-AD15ED52E61B.html> [3] Spine-Leaf Architecture: Introduction;

<https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series/switches/datasheet-c78-733553.html>

### QUESTION 3

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?

- A. By configuring IWA as an identity source
- B. By configuring LDAPS as an identity source
- C. By keeping the pre-configured defaults
- D. By replicating vSphere SSO configuration

**Correct Answer: D**

**Section:**

**Explanation:**

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.

Reference:

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/vspheresecurity-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.

Reference: VMware Cloud Foundation: Consolidated Architecture Deployment 4.0 on Dell EMC VxRail - Technical Overview (Page 24)

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vcf/techoverview/vmware-cloud-foundation-consolidated-architecture-dell-emc-vxrail.pdf>

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

### QUESTION 4

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

(Choose two.)

- A. SDDC Manager UI
- B. PowerCLI
- C. Cloud Builder UI
- D. vCenter UI
- E. REST API

**Correct Answer: A, E**

**Section:**

**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 1. 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 5

What is a valid procedure to replace an expired vSAN license in a VMware Cloud Foundation environment?

- A. 1 Add a new vSAN license to the SDDC Manager and vCenter Server.
- 2. Reassign the vSAN license to the cluster in the vCenter Server.
- 3. Remove the expired vSAN license from the SDDC Manager and vCenter Server.
- B. 1 Add a new vSAN license to the SDDC Manager.
- 2. Connect to SDDC Manager via SSH, and then restart Domain Manager using `systemctl restart domainmanager`.
- 3 Verify in the SDDC Manager whether a new vSAN license has been assigned to the cluster.
- C. 1 Add a new vSAN license to the vCenter Server.
- 2. Connect to SDDC Manager via SSH, and then restart Lifecycle Management using `systemctl restart lcm`.
- 3. Verify in the vCenter Server whether a new vSAN license has been assigned to the cluster.
- D. 1 Add a new vSAN license to the SDDC Manager.
- 2. Reassign the vSAN license to the cluster in the SDDC Manager.
- 3. Remove the expired vSAN license from the SDDC Manager

www.VCEplus.io

**Correct Answer: A**

**Section:**

**Explanation:**

<https://my-cloudy-world.com/2022/06/28/updating-a-vsan-license-in-vmware-cloud-foundation/>

#### QUESTION 6

What is a supported function of the vSphere Lifecycle Manager (LCM) in VMware Cloud Foundation (VCF)?

- A. Upgrade vCenter Server
- B. Upgrade SDDC Manager
- C. Downgrade VM hardware version
- D. Check hardware compatibility of hosts and clusters

**Correct Answer: A**

**Section:**

**Explanation:**

The VMware Cloud Foundation Specialist (v2) certification exam guide from VMware specifically mentions that "Upgrade vCenter Server to a new version using vSphere Lifecycle Manager (vLCM)" is a key knowledge area for the exam (source [1][2]: <https://www.vmware.com/educationservices/certification/vcf-specialist-exam.html>).

Additionally, search result [1] mentions that SDDC Manager is responsible for the life cycle management of vCenter Server and other components [3], indicating that vSphere LCM is specifically responsible for upgrading vCenter Server, a supported function. a supported function of vSphere Lifecycle Manager (LCM) in VMware Cloud Foundation (VCF) is Option A: Upgrade vCenter Server 1. vSphere Lifecycle Manager provides a framework for desired state lifecycle management of vSphere hosts in a cluster 2.

<https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-lifecycle-manager/GUID-F0635413-D52A-4115-A727-A1DCBDBAA292.html>

**QUESTION 7**

A systems administrator is tasked to deploy VMware Cloud Foundation (VCF) and has already deployed the VMware Cloud Builder appliance.

What is the next step the systems administrator should take?

- A. Deploy the Management Domain using VMware Cloud Builder.
- B. Prepare and configure the ESXi hosts to be used in the deployment.
- C. Configure the Management Domain using custom certificates.
- D. Upload the Deployment Parameter Workbook.

**Correct Answer: D**

**Section:**

**Explanation:**

According to the VMware Cloud Foundation documentation, after deploying the VMware Cloud Builder appliance, the next step is to create a deployment parameter workbook that defines the parameters required for the deployment. This workbook can be created manually or generated using the VMware Cloud Foundation Deployment Parameter Workbook Generator tool. Once the workbook is created, it must be uploaded to the VMware Cloud Builder appliance before deploying the Management Domain [1].

Therefore, the correct sequence of steps would be:

Deploy the VMware Cloud Builder appliance

Create or generate the Deployment Parameter Workbook

Upload the Deployment Parameter Workbook to the VMware Cloud Builder appliance Deploy the Management Domain using VMware Cloud Builder.

Reference: [1] VMware Cloud Foundation Planning and Preparation Guide -

<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.3/vcf-planning-preparation/GUIDDE823E7F-8682-4BFA-A293-91EFC48E3CAB.html>

**QUESTION 8**

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. Wols on ISCSI
- B. NFS
- C. vSAN and NFS
- D. vVols on NFS
- E. vSAN

**Correct Answer: B, E**

**Section:**

**Explanation:**

According to the VMware Cloud Foundation documentation on network requirements, vSAN requires only the vMotion and NFS networks to be configured in the network pool. This is because vSAN traffic can be carried over the vMotion network, and the NFS network is needed to support the use of NFS datastores [1].

NFS is a file-based storage protocol that can be accessed over IP networks. It does not require any special hardware or software, and can be accessed by any device that supports the NFS protocol [2].

As such, it only requires the NFS network to be configured in the network pool.

Therefore, the correct answers are B. NFS and E. vSAN.

**QUESTION 9**

An architect is tasked with deploying a new VI Workload Domain cluster to support the HR system.

The default storage policy must satisfy the following requirements:

- Support two host failures
- Use the least amount of hosts
- Maximize user capacity

Which configuration will satisfy these requirements?

- A. 4 Hosts, FFT=2, RAID 5/6

- B. 5 Hosts, FTT=2, RAID 1
- C. 4 Hosts, FTT=2, RAID 1
- D. 5 Hosts, FTT=2, RAID 5/6

**Correct Answer: A**

**Section:**

**Explanation:**

The requirements are to support two host failures, use the least amount of hosts, and maximize user capacity. RAID 5/6 and RAID 1 are both suitable options for a storage policy that can support two host failures. However, RAID 5/6 will be more efficient in terms of space utilization, allowing for a larger user capacity. In addition, the use of four hosts instead of five will minimize the number of hosts required, reducing costs associated with hardware and licensing.

According to VMware vSAN Design and sizing guide, the Failure to Tolerate (FTT) policy defines the number of replicas that must be kept for each object. In this case, FTT=2 means that two replicas of each object will be created. FTT=2 provides protection against the failure of two hosts. The Failure to Tolerate for the capacity tier (FFT) policy defines the number of parity fragments that must be kept for each stripe. RAID 5/6 requires FFT=1, which means that one parity fragment will be created for each stripe.

Therefore, the best answer is A. 4 Hosts, FTT=2, RAID 5/6.

<https://vmc.techzone.vmware.com/vmc-arch/docs/storage/vmc-aws-vsan-architecture> support two host failures (FTT=2), a minimum of four hosts is required. RAID 5/6 erasure coding can be used to maximize user capacity while using the least amount of hosts. This configuration will satisfy all three requirements: supporting two host failures, using the least amount of hosts, and maximizing user capacity.

#### QUESTION 10

Which service is integrated with VMware Cloud Foundation and enables a centralized and simplified lifecycle management of ESXi host?

- A. vRealize Suite Lifecycle Manager
- B. vSphere Lifecycle Manager
- C. Solutions Manager
- D. vCenter Lifecycle Manager

**Correct Answer: B**

**Section:**

**Explanation:**

The service that is integrated with VMware Cloud Foundation and enables a centralized and simplified lifecycle management of ESXi host is Option B: vSphere Lifecycle Manager (vLCM). vLCM enables you to create cluster images for centralized and simplified lifecycle management of ESXi hosts including firmware. When a VI workload domain cluster is created with an image, you can update and upgrade the ESXi version on all hosts in the cluster collectively. vSphere Lifecycle Manager (vLCM) is a key component of VMware Cloud Foundation (VCF) that enables centralized and simplified lifecycle management of ESXi hosts. It provides a single interface to manage host baselines, firmware and driver updates, and upgrades. With vLCM, administrators can create custom images for ESXi hosts, define baselines for host compliance, and apply updates to hosts in a coordinated manner. This helps to ensure consistency across the environment and reduce the risk of configuration drift. (source: VMware Cloud Foundation 4.x Architecture and Deployment Guide)

#### QUESTION 11

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)

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

**Correct Answer: C, D**

**Section:**

**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."

Reference: [1] Tanzu Kubernetes Cluster or Supervisor Cluster [1]: Which do I choose? -



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

#### QUESTION 12

Which two roles are provided by a local NSX Manager appliance? (Choose two.)

- A. Compliance
- B. Policy
- C. Controller
- D. Orchestrator
- E. Authorization

**Correct Answer: B, C**

**Section:**

**Explanation:**

The NSX Manager is a standalone appliance that hosts the API services, the management plane, control plane, and policy management. As a result of this combined format, you no longer need to install the manager and controllers as separate VMs. The NSX Manager has three built-in roles: policy, manager, and controller 1.

B. Policy: NSX Manager is responsible for the creation and management of NSX-T policies, which are used to define networking and security configurations and rules.

C. Controller: NSX Manager also serves as a central management point for NSX-T controllers, which are responsible for implementing and enforcing networking and security policies across the NSX-T environment.

Reference:

VMware Cloud Foundation Specialist (v2) Exam Guide, section 2.2

NSX-T Data Center Administration Guide, section "NSX Manager and NSX-T Controllers"

www.VCEplus.io

#### QUESTION 13

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

- A. VMware vSphere Evaluation
- B. VMware vSphere Standard
- C. VMware vSphere Enterprise Plus
- D. VMware Tanzu Basic

**Correct Answer: D**

**Section:**

**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.

Reference: VMware Cloud Foundation Specialist (v2) Exam Guide, Section 3: VMware vSphere with Tanzu, VMware Documentation: VMware Tanzu Kubernetes Grid Documentation 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.

Reference: VMware Cloud Foundation Specialist (v2) Exam Guide, Section 3: VMware vSphere with Tanzu, VMware Documentation: VMware Tanzu Kubernetes Grid Documentation

<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.2/com.vmware.vcf.vxrail.admin.doc/GUID-E8D0A432-8573-4DF5-9330-A4FE15F74128.html>

#### QUESTION 14

A systems administrator is tasked with creating a new VI workload domain that will leverage either an external NFS or a VMFS on FC storage as the principal storage.

Which action is required to fulfill this requirement?

- A. Create a new network pool for vMotion and vSAN networks.

- B. Create a new network pool only for vMotion network.
- C. Create a new network pool for vMotion, vSAN, and NFS networks.
- D. Create a new network pool for vMotion and NFS networks.

**Correct Answer: D**

**Section:**

**Explanation:**

This is because a network pool is a collection of subnets within an layer-2 network domain that includes information about subnets reserved for the vMotion and NFS networks that are required for adding a host to the SDDC Manager inventory<sup>2</sup>. The other options are not correct because they either include vSAN network, which is not needed for external NFS or VMFS on FC storage<sup>2</sup>, or they do not include NFS network, which is needed for external NFS storage<sup>1</sup>.

<https://infohub.delltechnologies.com/l/dell-storage-with-vmware-cloud-foundation-1/vmwarecloud-foundation-network-pool-configuration>

#### QUESTION 15

An architect is designing networking for a developer-ready infrastructure on VMware Cloud Foundation. During the discussion with the network team, a question comes up about the use of a routable CIDR range. Which item uses this type of range?

- A. ClusterIP
- B. vSphere Pod
- C. Ingress
- D. Kubernetes services

**Correct Answer: C**

**Section:**

**Explanation:**

This is because an ingress is a Kubernetes resource that exposes HTTP and HTTPS routes from outside the cluster to services within the cluster<sup>1</sup>. An ingress can use a routable CIDR range to assign IP addresses to the ingress controllers that handle the traffic routing.

#### QUESTION 16

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

**Correct Answer: A**

**Section:**

**Explanation:**

<https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-9A190942-BDB1-4A19-BA09-728820A716F2.html> To enable Workload Management (vSphere with Tanzu) on an existing workload domain cluster, a license for vSphere with Tanzu with sufficient CPU capacity must be added to the SDDC Manager inventory, and then assigned to the cluster in SDDC Manager. This is because vSphere Enterprise Plus license does not include an entitlement for vSphere with Tanzu. Therefore, Option B and Option C are incorrect. Option D is also incorrect since adding the license to both the SDDC Manager and vCenter Server is not necessary to enable Workload Management.

Reference: <https://docs.vmware.com/en/VMware-Cloud-Foundation/4.3/vcf-kubernetes/GUID-38B824A8-8C02-4223-BA3F-6AA2EDBF002A.html>

#### QUESTION 17

A systems administrator needs to apply a custom ESXi image to a host using VMware Imaging Appliance (VIA). Which statement is correct when preparing a host for imaging?

- A. Onboard NICs should be enabled on the server.
- B. VIA service does not support UEFI boot mode.
- C. VMware Cloud Builder appliance must be deployed in a tagged VLAN/Network.
- D. PXE Boot must be configured as the second boot option.



**Correct Answer: A**

**Section:**

**Explanation:**

This is because VIA service uses PXE boot to install ESXi on the servers, and it requires onboard NICs to be enabled and connected to an untagged VLAN/Network1.

According to VMware documentation on VMware Imaging Appliance, when preparing a host for imaging using VIA, it is recommended to enable the onboard NICs on the server. This enables the network adapter to participate in the boot sequence of the host to retrieve the image from the Imaging Appliance.

Here is the relevant quote from the documentation:

"To prepare the host, ensure that the onboard NICs are enabled on the server. During boot up, the server firmware detects the network adapter and adds it to the boot sequence list so that it can participate in network boot."

<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-deploy/GUID-735928E5-1DD7-44E5-BE32-E598230326AD.html>

**QUESTION 18**

An administrator is tasked with changing the password of the SDDC Manager super user account in a newly installed VCF environment.

Which method must the administrator use to complete this task?

- A. 1 Log in to SDDC manager UI as a user with the ADMIN role.
- 2. Go to Administration > Security > Password Management.
- 3. Select the SDDC Manager account from the component drop-down menu.
- 4. Click Rotate Now button.
- B. 1 SSH in to the SDDC Manager VM using the vcf user account.
- 2. Switch to the root user
- 3. Enter the passwd admin command.
- 4. Enter and retype the new password.
- C. 1 Log in to the SDDC manager UI as a user with the ADMIN role
- 2. Go to Developer Center > API Explorer
- 3. Expand APIs for managing users.
- 4. Update password for root user.
- D. 1 SSH in to the SDDC Manager VM using the vcf user account.
- 2. Switch to the root user.
- 3. Enter the passwd vcf command.
- 4. Enter and retype the new password.

**Correct Answer: A**

**Section:**

**Explanation:**

To change the password of the SDDC Manager super user account, the administrator should follow these steps:

Log in to the SDDC Manager UI as a user with the ADMIN role.

www.VCEplus.io

Go to Administration > Security > Password Management.

Select the SDDC Manager account from the component drop-down menu.

Click Rotate Now button.

This is the recommended method for changing the password of the SDDC Manager super user account, as documented in the VMware Cloud Foundation Administration Guide:

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

#### QUESTION 19

Which two requirements are needed to add new hosts to an existing VI workload in a VMware Cloud Foundation environment? (Choose two.)

- A. The host uses the same storage type as the existing cluster hosts.
- B. The host uses heterogenous hardware.
- C. The host uses a minimum of four network ports.
- D. The host uses the same network pool
- E. The host uses only the VLAN network.

**Correct Answer: A, D**

**Section:**

**Explanation:**

When adding new hosts to an existing VI workload domain in VMware Cloud Foundation, the new hosts must meet the following requirements:

A. The host uses the same storage type as the existing cluster hosts. D. The host uses the same network pool.

Reference:

<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.3/install-and-configure-vmware-cloudfoundation/GUID-7788F002-2FEA-426E-A153-7A78A1B200A7.html>The host must use the same storage type (vSAN or NFS) as other hosts in the cluster1.

The host must use a network pool that is compatible with other hosts in the cluster1.

The host must have at least two network ports for management traffic and two network ports for vSAN traffic (if using vSAN storage)2.

#### QUESTION 20

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

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

**Correct Answer: B, C**

**Section:**

**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 cluster12. These options are only available when using vSphere Lifecycle Manager images, not when using vSphere Lifecycle Manager baselines2.

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere-lifecyclemanager.doc/GUID-9A112231-AD7C-4EF5-AB6A-A8DAA704D307.html>

#### QUESTION 21

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

- A. vSAN, host overlay and vMotion network stretched across both sites
- B. DHCP on the NSX Edge overlay network
- C. DHCP on the host overlay network

- D. One witness host per site
- E. One witness host per vSAN stretched cluster
- F. vSAN: host overlay and vMotion network per data site

**Correct Answer: A, C, E**

**Section:**

**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-B38F02F74D40.html>

#### QUESTION 22

A VMware Cloud Foundation administrator has been tasked with replacing self-signed certificates with those signed by a third-party Certificate Authority. A security policy disallows the integration and use of Microsoft Active Directory Certificate Services and prefers an external provider.

Which two steps must be taken in order to configure these certificates? (Choose two.)

- A. Generate Certificate Signing Requests from SDDC Manager.
- B. Ensure that the external provider has Administrator rights in vCenter.
- C. Create and package the certificates in a domain\_name.tar.gz file
- D. Use the sddcmanager-ssl-util.sh utility to list and delete existing certificates.
- E. Generate public-private key pairs using the external provider.

**Correct Answer: A, C**

**Section:**

**Explanation:**

A. Generate Certificate Signing Requests from SDDC Manager - In order to replace the self-signed certificates with third-party signed certificates, the Certificate Signing Requests (CSRs) need to be generated. This can be done from the SDDC Manager UI.

C. Create and package the certificates in a domain\_name.tar.gz file - After the CSRs are generated, they can be used to obtain third-party signed certificates from a certificate authority. Once the certificates are obtained, they need to be packaged in a domain\_name.tar.gz file and uploaded to SDDC Manager.

Generate Certificate Signing Requests (CSRs) from SDDC Manager for each component that requires a certificate<sup>1</sup>. You can do this from the SDDC Manager UI or using an API call<sup>1</sup>.

Have the CSRs signed by a third-party Certificate Authority of your choice<sup>1</sup>. You can use any external provider that meets your security policy requirements.

Create and package the certificates in a domain\_name.tar.gz file according to the naming convention and folder structure specified by SDDC Manager<sup>1</sup>. You can use any compression tool that supports gzip format.

Upload and install the certificates using SDDC Manager UI or API<sup>1</sup>. You can also verify and troubleshoot the certificate installation using SDDC Manager.

<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-admin/GUID-80431626-B9CD-4F21-B681-A8F5024D2375.html>

#### QUESTION 23

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

- A. 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.

B. 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.

C. 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

D. 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.

**Correct Answer: A**

**Section:**

**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.

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.

Reference: VMware Cloud Foundation Specialist (v2) documentation on vmware.com (in particular, the "Replace Hosts for vSAN Stretched Cluster" section of the "vSAN" chapter).

#### QUESTION 24

What is required as part of enabling the Harbor Image Registry?

A. Storage Policy

B. Tanzu Enabled Cluster

C. Access Control

D. Resource Limits

**Correct Answer: A**

**Section:**

**Explanation:**

This is because according to Dell documentation<sup>1</sup>, to enable the Harbor Image Registry, you need to select the VM Storage Policy that will be used to store the images.

As part of enabling the Harbor Image Registry in VMware Cloud Foundation, a storage policy needs to be defined to specify the storage requirements for the registry. The storage policy should define the storage characteristics for the datastores where the registry will be deployed, including the redundancy level, disk type, and disk space. This is documented in the VMware documentation titled "Enabling Harbor Image

Registry in Workload Domains."

#### QUESTION 25

An administrator is tasked with enabling workload management for a VMware Cloud Foundation Management Workload Domain. This set of requirements was collected during the design workshops:

- Developers should be able to utilize vSphere Pods feature.
- Embedded harbor registry feature should be supported.
- Developers need to utilize persistent volumes across multiple provisioned vSphere Pods.

Which three actions will meet the requirements for this deployment? (Choose three.)

- A. Configure HA Proxy.
- B. Enable vSAN File Services.
- C. Enable vSphere HA and DRS in fully-automated mode
- D. Enable vSphere HA and DRS in partially-automated mode.
- E. Configure NSX Advanced Load Balancer.
- F. Configure NSX-T Networking.

**Correct Answer: B, C, F**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>2</sup>, these are some of the prerequisites for enabling workload management for a VMware Cloud Foundation Management Workload Domain:

You must have a vSphere cluster with NSX-T networking configured.

You must have vSAN File Services enabled on your cluster.

You must have vSphere HA and DRS enabled in fully automated mode on your cluster.

The other options are incorrect because they are not required or supported for this deployment.

<https://docs.vmware.com/en/VMware-Harbor-Registry/services/vmware-harbor-registry/GUIDindex.html>

#### QUESTION 26

Which component is upgraded when using the SDDC Manager management domain upgrade workflow in VMware Cloud Foundation?

- A. VMware Cloud Builder
- B. VMware vRealize Network Insight
- C. Workload Domain vCenter Server
- D. VMware NSX-T Manager nodes

**Correct Answer: A**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>1</sup>, the VMware Cloud Foundation Upgrade bundle upgrades the SDDC Manager appliance and Lifecycle Management, which are components of VMware Cloud Builder.

#### QUESTION 27

Which two features are supported when implementing NSX Federation? (Choose two.)

- A. DHCP dynamic binding
- B. Identity Firewall
- C. NAT operations
- D. Load Balancer
- E. DHCP

**Correct Answer: C, D**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>2</sup>, these are some of the features that are supported when implementing NSX Federation:

NAT operations: You can configure NAT rules for Tier-0 gateways across locations.

Load Balancer: You can configure load balancer services for Tier-0 gateways across locations.

#### QUESTION 28

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

**Correct Answer: A, D**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>3</sup>, 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 29

An administrator wants to scale an existing VMware Cloud Foundation (VCF) environment by adding a new VI Workload Domain. Which two storage options can be used for the new VI Workload Domain? (Choose two.)

- A. iSCSI principal storage with SMB supplemental storage
- B. vSAN principal storage with VMFS on FC supplemental storage
- C. vSAN principal storage with SMB supplemental storage
- D. vVOL principal storage with SCSI supplemental storage

**Correct Answer: B, D**

**Section:**

**Explanation:**

This is because according to VMware documentation, these are some of the storage options that can be used for a new VI Workload Domain: vSAN principal storage with VMFS on FC supplemental storage: You can use vSAN as your primary storage platform and add VMFS on FC as an additional storage option. vSAN principal storage with NFS supplemental storage: You can use vSAN as your primary storage platform and add NFS as an additional storage option.

#### QUESTION 30

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?

- A. 6
- B. 3
- C. 12
- D. 9

**Correct Answer: D**

**Section:**

**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.

Reference: VMware Cloud Foundation Specialist (v2) documentation on vmware.com (in particular, the "NSX-T Networking for VMware Cloud Foundation" chapter).

#### QUESTION 31

A vSphere administrator is tasked with enabling Workload Management on a VMware Cloud Foundation Workload Domain.

Which three components are configured as part of the Supervisor Cluster control plane after this task is completed? (Choose three.)

- A. Tanzu Kubernetes Grid Service
- B. kubectI-vSphere
- C. Kubernetes Grid Orchestrator
- D. Spherelet
- E. Kubernetes Mission Control
- F. Container Runtime Executive

**Correct Answer: A, C, D**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>2</sup>, these are some of the components that are configured as part of the Supervisor Cluster control plane after enabling Workload Management on a VMware Cloud Foundation Workload Domain:

Tanzu Kubernetes Grid Service: This service enables you to create and manage Tanzu Kubernetes clusters on vSphere with Tanzu.

Kubernetes Grid Orchestrator: This component manages the lifecycle of Tanzu Kubernetes clusters on vSphere with Tanzu.

Spherelet: This component runs on each ESXi host and acts as a kubelet agent that communicates with the Supervisor Cluster control plane.

#### QUESTION 32

Which action(s) can a developer perform on Kubernetes storage classes that are mapped from the VM Storage Policies?

- A. Access Only
- B. Access and Modify
- C. Access, Modify, and Delete
- D. Access, Create, and Delete

**Correct Answer: A**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>3</sup>, developers can only access Kubernetes storage classes that are mapped from VM Storage Policies. They cannot modify or delete them.

#### QUESTION 33

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?

- A. ESXi hosts need to be re-imaged and updated.
- B. ESXi hosts need to be decommissioned and re-imaged.
- C. ESXi hosts need to be re-imaged and rejoined.
- D. ESXi hosts need to be decommissioned and updated

**Correct Answer: B**

**Section:**

**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 34**

VCF design workshops were conducted, and the architect collected the following customer requirements for the newly planned VCF infrastructure:

- The new VCF infrastructure must target two zones: DEV/UAT and DMZ.
- The security team would like to have full management and network isolation between these two zones
- 12 hosts have been ordered for the solution.
- DEV/UAT workloads must comply with an erasure coding vSAN storage policy with the ability to tolerate the failure of two hosts.

Which workload domain sizing will be required to achieve these requirements?

- A. 12-hosts workload domain for both zones, having a 4-hosts DEV cluster a 4-hosts UAT cluster, and a 4-hosts DMZ cluster
- B. 12-hosts workload domain for both zones, having an 8-hosts DEV/UAT cluster, and a 4-hosts DMZ cluster
- C. 8-hosts DEV/UAT workload domain, having a 4-hosts DEV cluster and a 4-hosts UAT cluster, in addition to a 4-hosts DMZ workload domain, having a 4-hosts DMZ cluster
- D. 8-hosts DEV/UAT workload domain, having an 8-hosts DEV/UAT cluster, and a 4-hosts DMZ workload domain, having a 4-hosts DMZ cluster

**Correct Answer: D**

**Section:**

**Explanation:**

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/GUIDAD408FA8-5898-4541-9F82-FE72E6CD6227.html>

**QUESTION 35**

A VMware administrator is tasked to upgrade a VMware Cloud Foundation (VCF) environment that is running on Dell EMC PowerEdge servers.

During the ESXi software upgrade for the VI Workload Domain hosts, the administrator receives an error stating that the correct storage driver is not available, although the storage adapters are enabled in the BIOS.

Which action should the administrator take to fix this issue?

- A. Use the Dell EMC customized image for the ESXi build in the VCF bill of materials.
- B. Upgrade the storage adapter firmware to the latest version.
- C. Use the image for the ESXi build in the VCF bill of materials.
- D. Upgrade the BIOS firmware to the latest version.

**Correct Answer: A**

**Section:**

**Explanation:**

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/GUID-08911FD3-2462-4C1C-AE81-0D4DBC8F7990.html>

**QUESTION 36**

An architect needs to enable the workload management for a VCF VI workload domain and plans on using the minimum requirement for deploying a highly available NSX edge cluster.

What is the minimum requirement for this deployment?

- A. A cluster of three large NSX edge node
- B. A cluster of one large NSX edge node
- C. A cluster of two large NSX edge nodes
- D. A cluster of two small NSX edge nodes

**Correct Answer: C**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>1</sup>, a VI workload domain consists of one or more vSphere clusters that run customer workloads. To enable workload management for a VI workload domain, you need to have an NSX-T edge cluster that provides north-south connectivity and load balancing for Kubernetes clusters<sup>2</sup>. The minimum requirement for creating an NSX-T edge cluster is two edge nodes<sup>3</sup>. The size of the edge nodes depends on the throughput and features required, but for high availability, large size is recommended<sup>5</sup>.

<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-getting-started/GUID-C68FD810-D270-43F2-AEBF-D522BA1F402B.html>

**QUESTION 37**

A systems administrator wants to integrate Microsoft Certificate Authority with SDDC Manager and has already established a connection between the components.

Which pre-requisite step is required for this integration to work?

- A. Verify that the self-signed certificates have been replaced with signed certificates from Microsoft Certificate Authority.
- B. Verify that a customized certificate template has been configured on the SDDC Manager
- C. Verify that the Microsoft Certificate Authority Server has the RBAC roles configured on the same machine where the Certificate Authority role is installed.
- D. Verify that the Microsoft Certificate Authority Server has been configured for basic authentication

**Correct Answer: D**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>6</sup>, when integrating Microsoft Certificate Authority with SDDC Manager, you need to configure basic authentication on the Microsoft Certificate Authority Server and provide valid credentials when establishing a connection from SDDC Manager.

#### QUESTION 38

An administrator is tasked with preparing hosts for the deployment of a new Workload Domain in a VMware Cloud Foundation environment. The ESXi hosts have HBA cards that require the installation of a separate and the most recent VMware Installation Bundles (VIBs).

Which ESXi imaging method should the administrator use?

- A. VMware Imaging Appliance
- B. VMware vSphere Auto Deploy
- C. Download and burn the base ESXi Installer ISO Image to a DVD
- D. VMware vSphere Lifecycle Manager Images

**Correct Answer: A**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>1</sup>, the VMware Imaging Appliance (VIA) is a tool that automates ESXi host imaging and configuration for VCF deployments. It supports custom ESXi images that include additional drivers or VIBs that are not part of the base ESXi image<sup>2</sup>. You can download ESXi software and VIBs from VMware or third-party sources and upload them to VIA for creating custom images<sup>3</sup>.

#### QUESTION 39

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?

- A. node management domain and 3-node VI workload domain in a VCF standard model, while leveraging resource pools
- B. 3-node management domain and 3-node VI workload domain in a VCF standard model
- C. node management/workload domain in a VCF consolidated model, while leveraging resource pools
- D. node management/workload domain in a VCF consolidated model

**Correct Answer: B**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>45</sup>, 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>45</sup>.

#### QUESTION 40

The architect of a multi-site VMware Cloud Foundation solution is tasked with ensuring that the prerequisites for vSAN data at rest encryption have been achieved. The existing design calls for use of the vSphere Native Key Provider. NSX-T is configured with Federation, and both sites benefit from a stretched TO and T1 network topology.

A new security policy requires the use of vSphere Virtual Machine encryption, in addition to the atrest encryption already configured. During a failover test from Site-A to Site-B using Site Recovery Manager, the virtual machines were unable to power-on.

How does the design need to be changed to support the new requirement?

- A. Use a third-party KMS solution that allows for key replication.

- B. Use a third-party KMS solution at each site.
- C. Ensure that a TPM 2.0 certified module is installed on all ESXi hosts at Site-B.
- D. Ensure that the Site Recovery Manager service account has Cryptographer ReadKeyServersInfo privileges.

**Correct Answer: A**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>1</sup>, vSphere Native Key Provider (NKP) is a simple key management solution that does not support key replication across sites or clusters. It also does not support vSphere Virtual Machine encryption which requires a third-party KMS solution<sup>2</sup>.

Therefore, to enable both vSAN data at rest encryption and vSphere Virtual Machine encryption in a multi-site VCF solution, you need to use a third-party KMS solution that allows for key replication across sites.

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-54B9FBA2-FDB1-400B-A6AE-81BF3AC9DF97.html>

#### QUESTION 41

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?

- A. CNS-CSI
- B. Hostd
- C. Spherelet
- D. SPBM

**Correct Answer: D**

**Section:**

**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>.

#### QUESTION 42

An administrator is tasked with deploying a new VI Workload Domain into an existing VMware Cloud Foundation environment. Which three initial shared storage types are supported? (Choose three.)

- A. vSAN
- B. NFSv3
- C. SMB 3.0
- D. vVols
- E. VMFS on iSCSI
- F. NFSV4.1

**Correct Answer: A, B, E**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>6</sup>, these are three initial shared storage types that are supported for deploying a new VI Workload Domain into an existing VCF environment. You can also add other supported storage types after deploying the VI Workload Domain.

#### QUESTION 43

An administrator needs to upgrade the current VMware Cloud Foundation (VCF) environment from version 4.1 to 4.3, knowing that the environment has direct access to the internet.

Which steps should be performed to download the online bundles?

- A. 1 vSphere Lifecycle Manager checks depot.vmware.com.
- 2. The administrator accesses vSphere Lifecycle Manager.
- 3. The administrator downloads the bundles from vSphere Lifecycle Manager.

- B. 1. SDDC Manager checks depot.vmware.com.  
2. The administrator accesses vSphere Lifecycle Manager  
3. The administrator downloads the bundles from vSphere Lifecycle Manager
- C. 1. vSphere Lifecycle Manager checks depot vmware.com  
2. The administrator accesses SDDC Manager.  
3. The administrator downloads the bundles from SDDC Manager.
- D. 1. SDDC Manager checks depot.vmware.com.  
2. The administrator accesses SDDC Manager.  
3. The administrator downloads the bundles from SDDC Manager.

**Correct Answer: D**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>1</sup>, SDDC Manager is responsible for checking and downloading the online bundles for VCF upgrades from depot.vmware.com. The administrator can access SDDC Manager and download the bundles from Inventory > Workload Domains > MGMT> Update/Patches.

#### QUESTION 44

In which order should the VMware Cloud Foundation components in a Management Domain be upgraded?

- A. vRealize Suite, SDDC Manager and VMware Cloud Foundation services, NSX-T Datacenter, vCenter Server, ESXi servers  
B. SDDC Manager and VMware Cloud Foundation services, vRealize Suite, NSX-T Datacenter, ESXi servers, vCenter Server  
C. SDDC Manager and VMware Cloud Foundation services, vRealize Suite, NSX-T Datacenter, vCenter Server, ESXi servers  
D. SDDC Manager and VMware Cloud Foundation services. NSX-T Datacenter, vCenter Server, ESXi servers, vRealize Suite

**Correct Answer: D**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>2</sup>, this is the recommended order of upgrading VCF components in a Management Domain. The order ensures that there are no compatibility issues or dependencies between different components.

#### QUESTION 45

Which two use cases would require the creation of multiple clusters in a workload domain? (Choose two.)

- A. Different application licensing models  
B. Different host CPU types  
C. Different virtual machine workload types  
D. Different memory requirements  
E. Bare-Metal vs. virtual workloads

**Correct Answer: B, C**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>3</sup>, these are two use cases that would require the creation of multiple clusters in a workload domain. Different host CPU types may affect EVC compatibility and performance across clusters<sup>4</sup>. Bare-Metal vs. virtual workloads may have different storage and network requirements that need separate clusters<sup>5</sup>.

**QUESTION 46**

What limits can an administrator configure on a vSphere Namespace for a VMware Cloud Foundation Developer-Ready Workload Domain?

- A. Tanzu Kubernetes Cluster master nodes
- B. Secrets
- C. Tanzu Kubernetes Cluster worker nodes
- D. Persistent Volumes

**Correct Answer: B, D**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>1</sup>, these are two of the limits that an administrator can configure on a vSphere Namespace for a VMware Cloud Foundation Developer-Ready Workload Domain. Secrets are used to store sensitive information such as passwords or tokens for Kubernetes workloads<sup>2</sup>. Persistent Volumes are used to provide persistent storage for Kubernetes workloads<sup>3</sup>.

**QUESTION 47**

An administrator successfully finished restoring a SDDC Manager and now needs to verify its operation. Which tool should the administrator use for this verification?

- A. vRealize Operation plugin tool for VCF
- B. Support and Serviceability (SoS) tool
- C. SDDC Manager GUI restore health tool
- D. Ruby vSphere Console VCF check tool

**Correct Answer: B**

**Section:**

**Explanation:**

This is because according to VMware documentation<sup>5</sup>, this is the tool that an administrator should use for verifying the operation of SDDC Manager after restoring it from a file-based backup. The SoS tool can run various tests and checks on SDDC Manager and its components to ensure their health and functionality.

**QUESTION 48**

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

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

**Correct Answer: A**

**Section:**

**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 49**

A VMware architect has been asked to design a VMware Cloud Foundation solution for an online gaming company. The Chief Information Officer (CIO) of the company has asked the architect to focus on these requirements:

- The environment should be optimized for maximum hardware utilization.
- The environment should be highly available.

Which method meets these requirements and is supported for vCenter Server with VMware Cloud Foundation?

- A. Storage level snapshots
- B. vSphere HA
- C. vSphere FT
- D. vCenter HA

**Correct Answer: D**



**Section:****Explanation:**

This is because according to VMware documentation, this is one of the methods that meets these requirements and is supported for vCenter Server with VMware Cloud Foundation. vCenter HA provides high availability by creating an active-passive cluster of three vCenter Server nodes (one active, one passive, one witness). It also optimizes hardware utilization by allowing resource sharing among different workload domains through Enhanced Linked Mode (ELM).

**QUESTION 50**

An architect is designing overlay networking and routing for two VMware Cloud Foundation VI workload domains. The following requirements must be met in the design:

- NSX-T Bare Metal Edge nodes are required to satisfy a requirement for extremely low latency.
- The costs must be kept to a minimum.

Which design meets these customer requirements?

- A. Create separate NSX-T Manager clusters for each workload domain, and then create a single NSX-T Bare Metal Edge cluster to be shared between NSX-T Manager clusters.
- B. Place the Bare Metal Edge nodes in the same rack as the management domain hosts, and then add the NSX-T Bare Metal Edge cluster to the management domain NSX-T Manager cluster.
- C. Share an NSX-T Manager cluster between both workload domains, and then add the NSX-T Bare Metal Edge cluster to the shared NSX-T Manager cluster.
- D. Create separate NSX-T Manager clusters for each workload domain, and then create separate NSX-T Bare Metal Edge clusters for each workload domain.

**Correct Answer: C**

**Section:****Explanation:**

This is because according to VMware documentation, this is one of the supported design options for deploying NSX-T Bare Metal Edge nodes for VMware Cloud Foundation VI workload domains. Sharing an NSX-T Manager cluster between both workload domains can reduce costs by minimizing the number of NSX-T licenses required. Adding the NSX-T Bare Metal Edge cluster to the shared NSX-T Manager cluster can satisfy a requirement for extremely low latency by providing direct access to physical network interfaces.

**QUESTION 51**

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?

- A. vSphere Client
- B. SDDC Manager
- C. vSphere Update Manager
- D. vSAN PowerCLI

**Correct Answer: B**

**Section:****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 52**

Which three options are valid use cases for multiple clusters in a workload domain? (Choose three.)

- A. Workload domain with multiple availability zones each containing clusters
- B. Workload domain with multiple clusters in one rack
- C. One cluster stretching two workload domains
- D. Two clusters forming one stretched workload domain
- E. One stretched cluster across two regions and two workload domains
- F. One cluster distributed across two racks and one workload domain

**Correct Answer: A, B, D**

**Section:****Explanation:**

These are valid use cases for multiple clusters in a workload domain according to VMware documentation. A workload domain with multiple availability zones each containing clusters can provide high availability and

disaster recovery by distributing workloads across different physical locations within a region. A workload domain with multiple clusters in one rack can provide scalability and flexibility by allowing different configurations and policies for different clusters within a single logical unit. Two clusters forming one stretched workload domain can provide resiliency and fault tolerance by synchronizing data across two sites within a region.

#### QUESTION 53

A VCF administrator is preparing to configure scheduled backups for the SDDC Manager. What must the administrator register as an external component to complete this task?

- A. NFS server
- B. SFTP server
- C. iSCSI server
- D. SMB server

**Correct Answer: B**

**Section:**

**Explanation:**

This is because according to VMware documentation, this is what an administrator must register as an external component to complete this task of configuring scheduled backups for the SDDC Manager. SFTP server is one of the supported backup targets for SDDC Manager backups. The administrator can register an SFTP server by navigating to the SDDC Manager UI > Administration > Backup > Site Settings and clicking Register External.

#### QUESTION 54

An administrator is planning to deploy an edge cluster in a VMware Cloud Foundation environment.

Which three NSX components are automated during this deployment? (Choose three.)

- A. Segments for VM workloads
- B. Tier-0 VRF gateway configuration
- C. Tier-0 and Tier-1 gateway configuration
- D. Edge Uplink Profile configuration
- E. Edge VM deployment
- F. Transport Node Profile configuration

**Correct Answer: C, E, F**

**Section:**

**Explanation:**

These are NSX components that are automated during this deployment of an edge cluster in a VMware Cloud Foundation environment according to VMware documentation. Tier-0 and Tier-1 gateway configuration is automated by creating a default Tier-0 gateway with two uplinks and a default Tier-1 gateway with one downlink when deploying an edge cluster. Edge VM deployment is automated by deploying two edge VMs per edge node when deploying an edge cluster. Transport Node Profile configuration is automated by creating a transport node profile with N-VDS settings when deploying an edge cluster.

#### QUESTION 55

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

**Correct Answer: A**

**Section:**

**Explanation:**

This is because according to VMware documentation, this is the ideal design of VMware Cloud Foundation domains that the architect should recommend based on the Bill of Material for different vSAN ready nodes. This design can provide optimal performance and availability by using vSAN ready nodes with different disk group configurations for different domains based on their workload characteristics and requirements. The

management domain can use vSAN ready nodes with one disk group per node to provide sufficient capacity and redundancy for management components such as vCenter Server and NSX Manager. The workload domains can use vSAN ready nodes with two disk groups per node to provide higher performance and throughput for VI workloads such as VMs and containers.

#### QUESTION 56

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.

**Correct Answer: C**

**Section:**

**Explanation:**

According to VMware Cloud Foundation Specialist (v2) Exam1, one of the objectives is to “Describe how to expand an existing NSX-T Edge cluster”. The exam guide also states that “SDDC Manager provides a single point of management for password rotation” and that “password rotation must be performed through SDDC Manager”.

#### QUESTION 57

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.

**Correct Answer: B**

**Section:**

**Explanation:**

According to NSX Manager VM and Host Transport Node System Requirements<sup>23</sup>, an NSX-T management cluster formed using a medium-sized appliance can support up to 128 hypervisors.

Since each VI Workload Domain has 36 hosts (18 x 2), and there are three VI Workload Domains, the total number of hosts is 108 (36 x 3), which is within the limit of a medium-sized NSX-T Manager cluster.

#### QUESTION 58

An architect needs to create a VMware Cloud Foundation (VCF) VI Workload Domain design with these requirements:

- Design blueprint needs to be repeatable for additional regions
- Multiple availability zones
- Seven nodes per availability zone to host the workloads
- vSAN storage will be used

What is the maximum accepted latency supported by vSAN between the two availability zones'?

- A. 10 ms
- B. 100 ms
- C. 150 ms
- D. 5 ms

**Correct Answer: D**

**Section:**

**Explanation:**

According to Networking Requirements for vSAN<sup>1</sup>, the maximum network latency between the two main sites for stretched clusters is 5 ms RTT (round-trip time). This means that the latency between any two nodes in different availability zones should not exceed 5 ms.

**QUESTION 59**

During the design phase for a greenfield VMware Cloud Foundation (VCF) deployment, the following design decisions have been agreed upon:

- Stretched Cluster needs to be deployed
  - Identity and Access Management for VMware Cloud Foundation needs to be deployed
- The Infrastructure Architect is working with the client to fill the Planning and Preparation Workbook. The Option for Stretched Cluster has been set to 'Include'. The 'Identity and Access Management' is displaying an error, and its Final Result is stating 'Excluded'.

The Option for Stretched Cluster has been set to 'Include'. The 'Identity and Access Management' is displaying an error, and its Final Result is stating 'Excluded'.

Which option should be enabled in the Planning and Preparation Workbook to address the issue?

- A. Apply Signed Certificates
- B. Consolidated Management Domain
- C. Clustered Workspace One Access
- D. NSX Routing for Management Domain

**Correct Answer: D**

**Section:**

**Explanation:**

According to About Identity and Access Management for VMware Cloud Foundation<sup>23</sup>, one of the prerequisites for deploying Identity and Access Management for VMware Cloud Foundation is to enable NSX Routing for Management Domain in SDDC Manager. This allows SDDC Manager to communicate with Active Directory servers across different networks.

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vsan-planning.doc/GUIDAFF133BC-F4B6-4753-815F-20D3D752D898.html>

**QUESTION 60**

Which two steps must be performed to create a vSphere with Tanzu namespace? (Choose two.)

- A. Deploy a vSphere Cluster
- B. Define resource limits
- C. Enable Harbor Image Registry
- D. Assign permissions
- E. Use a DNS-compliant name

**Correct Answer: B, E**

**Section:**

**Explanation:**

According to How to Create a vSphere with Tanzu Namespace<sup>1</sup> and Create and Configure a vSphere Namespace<sup>2</sup>, two of the steps required to create a vSphere with Tanzu namespace are:

Define resource limits: You can specify CPU, memory, and storage limits for each namespace to control how much resources are available for the workloads running in that namespace.

Use a DNS-compliant name: You must provide a unique name for each namespace that is compliant with DNS naming conventions.

**QUESTION 61**

A VMware Cloud Foundation (VCF) architect is presented with a customer's requirements for an architecture that needs to achieve:

- Network high availability across workloads in two data center locations.
  - Maintain the least administrative overhead when performing day 2 operations.
  - Decrease the RTO of both management plane and data plane when site-wide failure occurs
- Which VCF design consideration should the architect recommend?

- A. VCF with NSX-T Bridging
- B. VCF with NSX-T L2-VPN
- C. VCF with NSX-T Federation
- D. VCF with NSX-T Multi-Site

**Correct Answer: C**

**Section:**

**Explanation:**

According to VMware Cloud Foundation Architecture Poster, VCF with NSX-T Federation provides network high availability across workloads in two data center locations by synchronizing network configuration and state across sites. It also simplifies day 2 operations by providing centralized management and policy enforcement across sites. It also reduces RTO by enabling fast failover of network services between sites.

**QUESTION 62**

www.VCEplus.io

An administrator is tasked with deploying an additional cluster within VI Workload Domain which has been created with vSAN as a principal storage. Which types of principal storage may the administrator configure during this process?

- A. iSCSI, NFS, vSAN; NFS v4.1
- B. NFS v4.1, VMFS on FC, vSAN
- C. vSAN, iSCSI, SMB3.0
- D. NFS, vSAN, NFS v3, VMFS on FC

**Correct Answer: B**

**Section:**

**Explanation:**

According to VMware Cloud Foundation Planning and Preparation Workbook, when adding an additional cluster within VI Workload Domain that has been created with vSAN as principal storage, you can choose from three types of principal storage:

NFS v4.1: You can use NFS version 4.1 datastores as principal storage if you have external NFS servers available.

VMFS on FC: You can use VMFS datastores on Fibre Channel (FC) SAN devices as principal storage if you have external FC SAN devices available.

#### QUESTION 63

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

- A. It runs as a VIB on all Supervisor Cluster ESXi hosts configured with the vSphere Networking Stack.
- B. It enables an ESXi hypervisor to act as a Kubernetes master node.
- C. It enables an ESXi hypervisor to act as a Kubernetes worker node.
- D. It starts and monitors vSphere pods running on the workload domain cluster.
- E. It communicates with the vSphere with Tanzu embedded Harbor registry.

**Correct Answer: B, C**

**Section:**

**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 64

Which two health checks can be performed using the SoS tool? (Choose two.)

- A. `--password-health`
- B. `--credential-health`
- C. `--system-check`
- D. `--health-check`
- E. `--esxi-health`

**Correct Answer: B, D**

**Section:**

**Explanation:**

According to Using the SoS Utility on VMware Cloud Builder<sup>2</sup> and Validate the SDDC Manager State<sup>3</sup>, two of the health checks that can be performed using the SoS tool are:

`--credential-health`: This option checks whether all credentials stored in SDDC Manager are valid and can be used to access various components or services.

`--health-check`: This option checks whether all components or services managed by SDDC Manager are healthy and running properly.

#### QUESTION 65

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

- A. Passwords can be common only for appliance users.
- B. Passwords can be different per user.
- C. Passwords must be different per user.
- D. Passwords must be common across all users.



**Correct Answer: C**

**Section:**

**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 66**

Which two functionalities does a NSX Tier-0 Gateway provide to a vSphere with Tanzu deployment?

(Choose two.)

- A. Gateway for Segments
- B. Layer 2 Switching
- C. Connectivity to all Tier-1 Gateways
- D. Downlink Connections to Segments
- E. Connectivity to physical networks and routers

**Correct Answer: C, E**

**Section:**

**Explanation:**

According to About Architecture and Design for a vSphere with Tanzu Workload Domain4, two of the functionalities that a NSX Tier-0 Gateway provides to a vSphere with Tanzu deployment are:

Connectivity to all Tier-1 Gateways: A Tier-0 Gateway connects to one or more Tier-1 Gateways that provide routing services for each namespace in vSphere with Tanzu.

Connectivity to physical networks and routers: A Tier-0 Gateway connects to external networks via uplink interfaces that can use static routing or dynamic routing protocols such as BGP.

**QUESTION 67**

Which two configuration steps must a VMware Cloud Foundation administrator apply to achieve north/south connectivity while setting up an edge VM node for a workload domain from the SDDC Manager user interface? (Choose two.)

- A. ToR Switches VRFs
- B. OSPF Configuration
- C. BGP Configuration
- D. vSphere VDS Uplinks
- E. NSX VDS Uplinks

**Correct Answer: D, E**

**Section:**

**Explanation:**

According to Deployment Model for the NSX-T Edge Nodes for a Virtual Infrastructure Workload Domain1, an NSX-T Edge node is an appliance that provides centralized networking services such as load balancing, NAT, VPN, and physical network uplinks. To achieve north/south connectivity for a workload domain from the SDDC Manager user interface, you need to configure two types of uplinks: vSphere VDS Uplinks: These are used to connect the NSX-T Edge node to the vSphere Distributed Switch (VDS) that provides network connectivity for all ESXi hosts in the workload domain cluster.

NSX VDS Uplinks: These are used to connect the NSX-T Edge node to the external networks via physical network interfaces on the ESXi host where it runs.

**QUESTION 68**

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

**Correct Answer: A, E**

**Section:**

**Explanation:**



According to VMware Cloud Foundation Planning and Preparation Workbook, two of the configurations that are part of the VMware Cloud Builder validation process are:

License key: Validates format, validity, and expiry for ESX, vSAN, vCenter Server, NSX, vRealize Suite, and Log Insight license keys  
Passwords: Validates specified passwords  
Checks for minimum length, invalid characters, and format

#### QUESTION 69

Which statement is true regarding NSX Manager configuration in a VMware Cloud Foundation environment?

- A. NSX Managers can be deployed to different VLANs.
- B. The cluster virtual IP address is used for API and GUI access to NSX Managers.
- C. Traffic is load-balanced across all NSX Managers while using the virtual IP address.
- D. The cluster virtual IP address is attached to all NSX Managers.

**Correct Answer: B**

**Section:**

**Explanation:**

According to VMware Cloud Foundation Planning and Preparation Workbook, a statement that is true regarding NSX Manager configuration in a VMware Cloud Foundation environment is:

The cluster virtual IP address (VIP) address must be used for API and GUI access to NSX Managers

#### QUESTION 70

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?

- A. 2 vSAN Ready Nodes, 192 GB RAM per server, and 4 10GbE NICs
- B. 8 vSAN Ready Nodes, 256 GB RAM per server, and 2 10GbE NICs
- C. 4 vSAN Ready Nodes, 192 GB RAM per server, and 2 10GbE NICs
- D. 6 vSAN Ready Nodes, 256 GB RAM per server, and 4 10GbE NICs

**Correct Answer: C**

**Section:**

**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

Reference:

<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.2/vcf-42-getting-started/GUID-0A0D7E16-C4D8-4B05-8C23-F7F2FF12DE64.html>

www.VCEplus.io