killexams.com

Vmware

5V0-31.22

VMware Cloud Foundation Specialist (v2)

ORDER FULL VERSION

Question: 9

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

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

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

Answer: B

Explanation:

According to NSX Manager VM and Host Transport Node System Requirements23, 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: 10

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.

Answer: A,B,D

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.

References:

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

VMware vSAN Documentation:https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vsan-planning.doc/GUID-9B7C9685-64C5-49C2-8E3C-CC2E47AFBC6F.html

Question: 11

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 at-rest 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-
D. Ensure that the Site Recovery Manager service account has Cryptographer ReadKeyServerslnfo privileges.

Answer: A

Explanation:

This is because according to VMware documentation1, 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 solution2. 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.

Question: 12

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

Answer: B

Explanation:

This is because according to VMware documentation45, 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 instances6. The minimum requirement for creating a management domain or a VI workload domain is three hosts45.

Question: 13

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.

Answer: D

Explanation:

after deploying the VMware Cloud Builder appliance, the next step a systems administrator should take when deploying VMware Cloud Foundation (VCF) is Option D: Upload the Deployment Parameter Workbook 1. After uploading and validating the deployment parameter workbook, you can perform bring-up of the management domain 1.

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. Oncethe 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:

â Deploy the VMware Cloud Builder appliance

â Create or generate the Deployment Parameter Workbook

â Upload the Deployment Parameter Workbook to the VMware Cloud Builder appliance

â 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/GUID-DE823E7F-8682-4BFA-A293-91EFC48E3CAB.html

Question: 14

Remove the expired vSAN license from the SDDC Manager

- wrong

Answer: D

Explanation:

To replace an expired vSAN license in a VMware Cloud Foundation environment, follow these steps:

Add a new vSAN license to the SDDC Manager. This will make the new license available for use within the VCF environment.

Reassign the vSAN license to the cluster in the SDDC Manager. This step ensures that the new license is applied to the cluster, replacing the expired license.

Remove the expired vSAN license from the SDDC Manager. This step helps to avoid confusion and ensures that the expired license is no longer used in the environment.

Question: 15

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, Wols, vSAN; NFS v4 1
B. NFS v4.1, VMFSon FC, vSAN
C. vSAN, iSCSI, SMB3.0

D. Wols, vSAN, NFS v3, VMFS on FC

Answer: B

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:

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

â 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: 16

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

Answer: D

Explanation:

This is because according to VMware documentation34, 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 requirements5.

Question: 17

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

Answer: A,E,F

Explanation:

To deploy a stretched cluster in a VMware Cloud Foundation environment, the following components are required:

vSAN, host overlay, and vMotion network stretched across both sites: This is necessary to enable data and vMotion traffic between the two sites, ensuring that the stretched cluster can function properly.

One witness host per vSAN stretched cluster: The witness host acts as a tiebreaker and provides quorum for the vSAN stretched cluster. It is typically deployed at a third location to avoid split-brain scenarios.

vSAN: host overlay and vMotion network per data site: Each data site requires its own host overlay and vMotion network to ensure proper communication and functionality within the stretched cluster.

Question: 18

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

Answer: A,C,E

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:

â 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.

â 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: 19

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

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

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

different stretched Workload Domain cluster.

Answer: C

Explanation:

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

Question: 20

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.

Answer: C

Explanation:

When implementing stretched clusters in an environment with multiple Availability Zones (AZs), it is essential to have a highly available Layer 3 gateway for both the workload domain and Edge overlay networks across the AZs. This ensures seamless communication between the workloads and resources in the different AZs and provides the necessary resiliency and fault tolerance required in such a distributed architecture.