

VMware 3V0-23.25


Advanced VMware Cloud Foundation 9.0 Storage

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/3v0-23-25>

Latest Version: 6.2

Question: 1

An administrator has been tasked with recommending a principal storage type for a new cluster within an existing VMware Cloud Foundation (VCF) Workload Domain.

The following information has been provided:

- The customer has a 25GbE capable network.
- The customer has an existing 3rd party storage solution that supports connectivity by either FibreChannel and NFS
- The new cluster will consist of two host servers that will be recycled from a previous project.
- Each of the host servers have the following configuration:
 - o One 500GB Enterprise-grade SSD drive.
 - o Two 2-Port IOGbE network cards.
- There is no budget for additional hardware.

Which approach should the administrator recommend for the principal storage type?

- A. Fibre Channel
- B. NFS v4.1
- C. NFS v3
- D. Fibre Channel over Ethernet

Answer: C

Explanation:

NFS v3 is the best recommendation because the scenario provides only Ethernet-based host connectivity and no budget for additional hardware. Fibre Channel would require appropriate FC HBAs, SAN zoning, and array presentation to all hosts. The provided hosts only have Ethernet NICs and local boot/storage devices, so they are not prepared for FC without additional hardware. VMware Cloud Foundation supports NFS as principal storage for VI workload domains, but the supported principal NFS protocol version is NFS v3. NFS v4.1 is supported as supplemental storage, not as principal storage for this use case. The documentation also states that when NFS is used as principal storage with vSphere Lifecycle Manager images, only two hosts are required, matching this two-host scenario. Fibre Channel over Ethernet is not the appropriate supported principal storage option in this VCF storage model. Therefore, NFS v3 uses the existing network capability, avoids new hardware spend, and aligns with the supported VCF principal storage model. Reference topics: Principal Storage, NFS Storage Model, VI Workload Domain Storage, Fibre Channel Prerequisites.


Question: 2

DRAGDROP


A VMware Cloud Foundation (VCF) Workload Domain is requested to be deployed with the following information:

- . 6 blade style hosts with no local storage beyond the operating system.
- . 4 25 Gb networking cards installed in each host.
- . A 30 TB external array configured to support NVMe/TCP only.
- . 2 dVS switches, one configured for storage isolation and one for all other traffic.
- . NVMe/TCP multi-path configuration required.
- . Existing Management Domain is deployed with VCF.

Place the steps for importing VCF on to this configuration.



Answer:



Question: 3

A storage administrator is being presented with the following VMware Cloud Foundation (VCF) architectural details:

- The applications data will require 2.5 PB of capacity.
- The production applications will be hosting archival solutions and gateways.
- There will be some applications deployed for the purpose of testing and development.

What is the optimal principal storage that the administrator can recommend?

- A. vSAN ESA Storage Clusters
- B. vSAN ESA
- C. vSAN OSA - All Flash
- D. vSAN OSA - Hybrid

Answer: A

Explanation:

vSAN ESA Storage Clusters are the optimal choice because the workload profile is capacity-heavy and benefits from disaggregated storage. The requirement for 2.5 PB of application data, archival workloads, gateways, and development/test applications indicates a need to scale storage capacity independently from compute resources. VMware Cloud Foundation describes a vSAN storage cluster as a disaggregated storage solution based on vSAN ESA. It provides storage resources but not compute resources, and its datastore can be mounted by vSAN compute clusters or vSAN HCI clusters. This model is designed for storage-dense environments where capacity and performance can be expanded without adding unnecessary compute. A standard vSAN ESA HCI cluster is excellent for high-performance workloads, but it scales compute and storage together. vSAN OSA all-flash and hybrid use the older disk-group architecture and are less suitable for this large-capacity, storage-centric design. Therefore, the principal storage recommendation should be vSAN ESA Storage Clusters to meet scalability, performance, and capacity growth requirements. Reference topics: vSAN Storage Cluster, Disaggregated Storage, vSAN ESA, Storage Models.

Question: 4

A storage architect is designing a VMware Cloud Foundation (VCF) Workload Domain with the following requirements:

- vSAN File Services hosting multiple file shares.
- Each department requires distinct protection levels and placement rules.

Which option satisfies the requirements?

- A. Create multiple IP Pools and bind each to a file server to be able to assign different storage policies.
- B. Separate file shares by creating individual File Service clusters per department.
- C. Make sure the file share server VMS are tied to a vSphere Compute Policy during creation.
- D. Assign different vSAN Storage Policies to each file share during creation.

Answer: D

Explanation:

Assigning different vSAN Storage Policies to each file share during creation satisfies the requirement because vSAN File Service allows each file share to be created with its own storage policy. The Create File Share workflow includes a Storage Policy selection, which controls how

the file share's underlying vSAN objects are protected and placed. This is the correct mechanism when each department requires distinct protection levels or placement rules. IP pools and file server bindings provide network endpoint capacity and service placement, but they do not define the protection level of file share data. Creating separate File Service clusters per department is unnecessary and operationally inefficient because file shares can be differentiated by policy within the same File Service deployment. vSphere Compute Policies influence compute placement, not vSAN object protection or file-share resilience. Therefore, the department-specific requirement is fulfilled directly by assigning the appropriate vSAN Storage Policy to each file share at creation time. Reference topics: vSAN File Service, Create File Share, Storage Policy, Placement and Protection Rules.

Question: 5

An administrator is managing a local vSAN 3-node cluster running the Express Storage Architecture (ESA) as part of a VMware Cloud Foundation (VCF) Workload Domain.

The following parameters apply at the cluster level:

- HA is enabled
- Reserved Failover Capacity = 25% for both CPU and Memory

The administrator is tasked with configuring the vSAN storage policy so it will tolerate the maximum number of failures.

What two protection levels can the administrator configure? (Choose two.)

- A. RAID-1 (Mirroring) 3 Failures.
- B. RAID-1 (Mirroring) 2 Failures.
- C. No data redundancy.
- D. RAID-6 (Erasure Coding) 2 Failures.
- E. RAID-1 (Mirroring) 1 Failure.
- F. RAID-5 (Erasure Coding) 1 Failure.

Answer: E, F

Explanation:

In a three-node vSAN ESA cluster, the maximum practical storage policy failure tolerance is one failure. RAID-1 mirroring with Failures to tolerate set to 1 is supported because vSAN can place mirrored data components across the available hosts while maintaining quorum. RAID-5 erasure coding with Failures to tolerate set to 1 is also supported in vSAN ESA three-host clusters. The documentation specifically states that RAID-5 in vSAN ESA supports three ESX host clusters. RAID-6 with FTT=2 is not supported because it requires at least six ESX hosts in the vSphere cluster. RAID-1 with FTT=2 would also require additional fault domains because vSAN requires $2n+1$ fault domains for n failures tolerated. HA reserved failover capacity for CPU and memory does not change the number of vSAN fault domains available for storage object placement. It only affects compute admission control. Therefore, the two valid maximum protection levels for this three-node ESA cluster are RAID-1 mirroring with one failure and RAID-5 erasure coding

with one failure. Reference topics: vSAN ESA Storage Policies, Failures to Tolerate, RAID-1 Mirroring, RAID-5 Erasure Coding, RAID-6 Requirements.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>