

VMware HCI Master Specialist

VMware 5V0-21.21

Version Demo

Total Demo Questions: 7

Total Premium Questions: 73

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Architectures and Technologies	13
Topic 2, VMware Products and Solutions	10
Topic 3, Planning and Designing	19
Topic 4, Installing, Configuring, and Setup	7
Topic 5, Performance-tuning, Optimization, and Upgrades	6
Topic 6, Troubleshooting and Repairing	11
Topic 7, Administrative and Operational Tasks	7
Total	73

QUESTION NO: 1

An architect is working with vSAN and setting the fault domains to support FTT=1.

How many fault domains will be needed?

- A. 2
- B. 3
- C. 4
- D. 1

ANSWER: B

Explanation:

3 is the required number of fault domains to support a vSAN storage policy with failures to tolerate set to 1. FTT=1 means that vSAN must be able to sustain the loss of one complete fault domain while continuing to provide the required object availability. To do that, vSAN needs one fault domain for the primary data component, one separate fault domain for the replica component, and an additional fault domain that can fail. Therefore, the minimum number follows the rule of $2 \times \text{FTT} + 1$ fault domains; for FTT=1, this equals 3.

A fault domain is a logical grouping of hosts intended to represent a common failure boundary, such as a rack that shares power and network infrastructure. When fault domains are configured, vSAN places data replicas across distinct fault domains rather than merely across different hosts. This placement protects workloads from a rack-level or similarly scoped infrastructure failure, provided the fault-domain design accurately reflects the physical failure domains in the environment. VMware documents the fault-domain requirements and the relationship between failures to tolerate and available fault domains in the [vSAN fault domains documentation](#) and its [vSAN cluster design guidance](#).

QUESTION NO: 2

In a vSAN stretched cluster, which value must be set in the vSAN policy if there is no requirement for data mirroring across sites?

- A. SFTT = 0
- B. SFTT = 1
- C. PFTT = 1
- D. PFTT = 0

ANSWER: A

Explanation:

SFTT = 0 is required when a vSAN stretched-cluster workload does not need data mirrored between the preferred and secondary sites. Site failures to tolerate (SFTT) is the stretched-cluster policy rule that controls whether an object receives a replica in the opposite data site. A value of zero specifies that the object does not require site-level redundancy, so vSAN can keep the object's data components within one site rather than maintaining a mirrored copy across the intersite link.

This setting is appropriate for workloads whose availability requirements do not include surviving the loss of an entire data site, or where cross-site mirroring would add unnecessary capacity consumption and network traffic. It is independent of the object's local availability requirements: failures to tolerate (PFTT) continues to define protection against host, disk, or other local failures according to the selected failure-tolerance method. Consequently, SFTT must be explicitly set to zero for the no-cross-site-mirroring requirement, while PFTT should be chosen separately to meet the workload's local resilience target.

See the VMware vSAN stretched-cluster planning guidance in the [vSAN Planning documentation](#) and the [VxRail vSAN Stretched Cluster Planning Guide](#).

QUESTION NO: 3

A company has engaged a consultant to upgrade an existing vSAN cluster to vSAN 7.0 U1. The company wants to ensure that the same vSAN process can be used in the future.

During the discovery phase, the consultant found the following information about the existing environment:

- The vCenter Server is currently version 7.0.
- The vSAN Cluster has the following configuration:
 - vSAN version: 7.0
 - Number of vSAN nodes: 6
 - Encryption: enabled
 - Deduplication and Compression: enabled
 - Fault Domains: 1
 - vSAN Capacity Utilization: 60%
- Each vSAN node has the following configuration:
 - ESXi version: VMware vSphere 7.0 - CPU: 2 processors, 20 cores - RAM: 1024GB RAM.
 - Disk: 2 Cache SSDs and 6 Capacity SSDs
 - Network: 4 x 10GbE
- All current hardware (which is from a single vendor) is listed on the vSAN Compatibility Guide for vSAN 7.

Which three recommendations should the consultant make to ensure that the vSAN cluster upgrade is completed? (Choose three.)

- A.** Upgrade all vSAN nodes to VMware vSphere 7.0 U1 using the baselines capability within VMware Update Manager (VUM).
- B.** Disable vSAN Encryption before starting the upgrade process.
- C.** Upgrade all vSAN nodes to VMware vSphere 7.0 U1 using the images capability within VMware Update Manager (VUM).
- D.** Set VMware Distributed Resource Scheduling (DRS) to partially automated.
- E.** Upgrade to VMware vCenter Server 7.0 U1.
- F.** Choose the Allow reduced redundancy option.
- G.** Upgrade vCenter Server to 7.0 U1, upgrade all vSAN nodes to vSphere 7.0 U1 using VMware Update Manager baselines, and choose the Allow reduced redundancy option.

ANSWER: G

Explanation:

The appropriate combined recommendation is to first upgrade vCenter Server to version 7.0 Update 1, then remediate the vSAN hosts with VMware Update Manager baselines, while allowing reduced redundancy for the remediation operation. vCenter Server should be upgraded before the ESXi hosts so it can manage the target host version and provide the Update Manager functionality needed to orchestrate the cluster upgrade. A baseline-based workflow provides a repeatable process that can be retained and used for later maintenance cycles.

During sequential host remediation, a host enters maintenance mode and vSAN evaluates whether objects can remain available and compliant. Enabling the reduced-redundancy remediation setting permits the process to continue when full policy redundancy cannot be maintained temporarily during a host upgrade. This is particularly relevant when capacity and fault-domain constraints can make full evacuation or immediate object re-protection impractical. The consultant should still validate object health, available capacity, and the vSAN compatibility status before remediation, then monitor

resynchronization after each host returns to the cluster. VMware documents the required vCenter-before-host upgrade order and supported upgrade workflows in its [vSphere Upgrade documentation](#) and provides vSAN lifecycle guidance in the [vSAN documentation](#).

QUESTION NO: 4

An architect is evaluating vSAN Express Storage Architecture for a new workload cluster.

Which two characteristics are provided by vSAN ESA? (Choose two.)

- A. A single-tier NVMe storage architecture
- B. Native vSAN snapshots
- C. Disk groups containing one cache device and capacity devices
- D. A dedicated flash cache tier for every host
- E. Magnetic disks as the ESA capacity tier

ANSWER: A B

Explanation:

vSAN ESA uses a single-tier storage architecture in which NVMe devices contribute to both capacity and performance. This differs from vSAN OSA, which uses cache and capacity tiers organized in disk groups.

ESA also provides native snapshots designed to reduce snapshot-consolidation impact and improve snapshot performance compared with traditional redo-log based snapshots. Native snapshots are integrated with the ESA log-structured file system and vSAN object architecture.

ESA does not use disk groups, and it does not require a dedicated cache device for each group. Traditional magnetic capacity devices are also not part of the ESA storage design.

See the [vSAN Express Storage Architecture overview](#).

QUESTION NO: 5

An administrator plans to place a vSAN host into maintenance mode by using the No data migration option.

Which two statements are true? (Choose two.)

- A. vSAN object components remain on the host.
- B. Objects can become noncompliant while the host is in maintenance mode.
- C. vSAN evacuates all accessible components before the host enters maintenance mode.
- D. vSAN automatically creates a full replica of every object before maintenance begins.
- E. The option guarantees that all objects remain policy compliant.

ANSWER: A B

Explanation:

No data migration does not evacuate vSAN object components from the host. Components remain on the host while it is in maintenance mode, so the administrator must verify that affected objects remain accessible through their remaining components.

This option is appropriate only when the maintenance period is short and the selected host can be taken offline without violating the availability requirement of the workloads. It does not preserve full policy compliance because vSAN does not create replacement components before the host enters maintenance mode. If the host remains unavailable beyond the configured object repair delay, vSAN can begin rebuilding absent components elsewhere, provided sufficient resources are available.

See the [vSAN Administration documentation](#) for maintenance-mode behavior.

QUESTION NO: 6

An administrator is configuring vSAN File Services to provide SMB file shares to users.

Which directory service must be configured to support SMB share access?

- A. Active Directory
- B. VMware Directory Service
- C. vCenter Single Sign-On
- D. vSAN Skyline Health

ANSWER: A

Explanation:

Active Directory must be configured for SMB file shares provided by vSAN File Services. The file service uses the Active Directory domain to authenticate SMB users and groups and to apply share permissions. The administrator joins the vSAN File Services domain configuration to the appropriate Active Directory domain before creating SMB shares that require domain-based access control.

NFS shares can use export policies and, when required, Kerberos-based authentication, but SMB access relies on Active Directory integration. See the [vSAN File Services documentation](#).

QUESTION NO: 7

A vSAN administrator observes resynchronization activity in a cluster after a configuration change.

Which two events can directly trigger vSAN resynchronization? (Choose two.)

- A. A storage policy is changed from FTT=0 to FTT=1.
- B. A host remains unavailable beyond the object repair delay.
- C. A new host is added to the cluster.
- D. The vCenter Server inventory is reorganized.
- E. A virtual machine is renamed.

ANSWER: A B

Explanation:

A change that increases the Failures To Tolerate requirement causes vSAN to create additional object components so the object can satisfy the revised storage policy. Copying data to those new components generates resynchronization traffic until the object becomes compliant.

A host or disk failure can also trigger resynchronization after the configured object repair delay expires. vSAN rebuilds affected components on eligible remaining devices and hosts to restore the required availability and compliance. Adding a

host to a cluster does not, by itself, require existing objects to be resynchronized; vSAN does not automatically rebalance all data merely because additional capacity becomes available. See the [vSAN Monitoring and Troubleshooting documentation](#).