



# 5V0-22.23<sup>Q&As</sup>

VMware vSAN Specialist v2

**Pass VMware 5V0-22.23 Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/5v0-22-23.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by VMware  
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





### QUESTION 1

What is the purpose of host rebuild reserve in vSAN?

- A. Reserves space for internal operations
- B. Reserves space in case of single host failure
- C. Stores vSphere HA heartbeats
- D. Allocates capacity for vCLS

Correct Answer: B

Explanation: The host rebuild reserve is a feature that allows vSAN to reserve space in the cluster for vSAN to be able to repair in case of a single host failure. This reservation is set to one host worth of capacity, which means that if one host in the vSAN cluster fails and no longer contributes storage, there is still sufficient capacity remaining in the cluster to rebuild and re-protect all vSAN objects. This feature prevents the creation of new VMs or powering on VMs if such operations consume the reserved space. By default, the host rebuild reserve is disabled, but it can be enabled in the vSAN Services configuration. The other options are not related to the host rebuild reserve. References: vSAN Capacity Management in v7.0U1; Configure Reserved Capacity

---

### QUESTION 2

vSAN requires that the virtual machines deployed on the vSAN datastores are assigned at least one storage policy, but the administrator did not explicitly assign a storage policy when provisioning the new VM.

What is the result of this situation?

- A. The VM provisioning will fail.
- B. The VM objects will be protected based on the vSAN Default Storage Policy configurations.
- C. The vSphere Web Client will choose the last vSAN Storage Policy used.
- D. No data protection will be applied to the VM objects.

Correct Answer: B

Explanation: If the administrator did not explicitly assign a storage policy when provisioning a new VM on a vSAN datastore, the result is that the VM objects will be protected based on the vSAN Default Storage Policy configurations. The vSAN Default Storage Policy is assigned to all VM objects if no other vSAN policy is assigned when provisioning a VM. The default policy contains vSAN rule sets and a set of basic storage capabilities, such as Failures to tolerate set to 1, Number of disk stripes per object set to 1, and Thin provisioning. The other options are not correct. The VM provisioning will not fail, as vSAN requires that every VM has at least one storage policy. The vSphere Web Client will not choose the last vSAN Storage Policy used, as it will always apply the default policy if no other policy is selected. No data protection will not be applied to the VM objects, as they will have at least one replica based on the default policy. References: About the vSAN Default Storage Policy; Using vSAN Policies

---

### QUESTION 3



A vSAN administrator needs to enable vSAN ESA.

Which two requirements need to be met? (Choose two.)

- A. vSAN Build Your Own configuration
- B. vSAN Standard license
- C. vSAN Witness Appliance
- D. vSAN Advanced license
- E. vSAN ReadyNodes configuration

Correct Answer: BE

Explanation: To enable vSAN ESA, two requirements that need to be met are: vSAN Standard license or higher, and vSAN ReadyNodes configuration. vSAN Standard license or higher is required to use vSAN ESA, as it is a feature that is only available in vSAN 8.0 or later versions. vSAN ESA is an optional, alternative architecture to vSAN OSA that is designed to process and store data with higher efficiency, scalability, and performance. vSAN ReadyNodes configuration is required to use vSAN ESA, as it is a hardware configuration that is pre-configured, tested, and certified for VMware Hyper-Converged Infrastructure Software. Each vSAN ReadyNode is optimally configured for vSAN ESA with the required amount of CPU, memory, network, and storage NVMe devices. The other options are not correct. vSAN Build Your Own configuration is not supported for vSAN ESA, as it might not meet the hardware requirements or compatibility for vSAN ESA. vSAN Witness Appliance is not required to use vSAN ESA, as it is only needed for stretched cluster or two-node cluster configurations. References: vSAN Express Storage Architecture; vSAN ReadyNode Hardware Guidance

#### QUESTION 4

An administrator is performing maintenance on the hosts in a four-node vSAN cluster and has selected the "Ensure Accessibility" maintenance mode option. All VMs are running with the Default Storage Policy which has not been modified from the default settings.

While one of the hosts in the cluster is down for firmware upgrade, a second host suddenly loses network connectivity to the remaining hosts.

How will the cluster be affected?

- A. VMs might experience data loss
- B. Cluster will still be fully operational
- C. All VMs in the cluster will be inaccessible
- D. The backend performance metrics will be lost

Correct Answer: A

Explanation: If two hosts in a four-node vSAN cluster are down, the cluster might experience data loss because the default storage policy has a Primary level of failures to tolerate (PFTT) of 1, which means that vSAN can tolerate only one host failure. The Ensure accessibility maintenance mode option does not guarantee full data redundancy, but only ensures that all accessible VMs remain accessible. If another host fails while one host is in maintenance mode, some VMs might lose access to their data components and become unavailable or corrupted. References: vSAN Maintenance Mode Options; vSAN Cluster Configuration Limits



### QUESTION 5

An architect is designing a vSAN stretched cluster and needs to ensure that data remains on a given site in case of a network partition between the sites.

Which configuration would do this?

- A. Preferred and secondary sites
- B. vCenter High Availability
- C. Distributed Resource Scheduler
- D. IvSoh ere High Availability

Correct Answer: A

Explanation: In a vSAN stretched cluster configuration, both data sites are active sites, but one site must be designated as the preferred site and the other site as the secondary or nonpreferred site. This configuration helps to ensure that data remains on a given site in case of a network partition between the sites. If the network connection between the two active sites is lost, vSAN continues operation with the preferred site, unless it is resyncing or has another issue. The site that leads to maximum data availability is the one that remains in operation. The other options are not relevant to this scenario. References: Introduction to Stretched Clusters; vSAN Stretched Cluster Guide

---

### QUESTION 6

A vSAN administrator has a group of requirements from the application team, which mandates spreading the components across storage devices as much as possible.

What should the vSAN Administrator consider to achieve such a requirement for building a new vSAN cluster? (Choose two.)

- A. Configure disk striping in OSA
- B. Configure disk striping in ESA
- C. Enable Force Provisioning in OSA
- D. Enable deduplication for vSAN
- E. Create a dedicated Storage Pool in ESA

Correct Answer: AD

Explanation: To spread the components across storage devices as much as possible, the vSAN administrator can configure disk striping in either OSA or ESA. Disk striping is a policy attribute that defines the number of capacity devices across which each replica of a storage object is striped. A higher number of stripes can result in better performance and availability, but also consumes more storage space. Disk striping can be configured in OSA by using the Number of disk stripes per object policy attribute, or in ESA by using the Striping Width policy attribute<sup>12</sup>  
References: 1: VMware vSAN Specialist v2 Exam Preparation Guide, page 14 2: VMware vSAN Design and Sizing Guide, page 32

---

**QUESTION 7**

A six-node vSAN ESA cluster contains multiple virtual machines, and a vSAN storage policy with the rule "Failures to tolerate" set to "1 failure - RAID-5 (Erasure Coding)" is assigned. A vSAN administrator has changed the rule in the assigned policy to "2 failures - RAID-6 (Erasure Coding)".

What is the result of this change?

- A. No changes occur until the policy is reapplied.
- B. The changes are queued for 60 minutes.
- C. The policy change is rejected immediately.
- D. The updated policy is serially applied to the virtual machines.

Correct Answer: D

Explanation: The updated policy is serially applied to the virtual machines is the correct answer because changing the rule in the assigned policy will trigger a policy compliance check and a resynchronization of the affected objects. The policy change will not be rejected, queued, or ignored, as it is a valid and supported operation. However, the policy change will not be applied in parallel, as that would cause too much network and disk traffic. Instead, the policy change will be applied one virtual machine at a time, starting with the most critical ones, until all virtual machines are compliant with the new policy. References: VMware vSAN Specialist v2 Exam Preparation Guide, page 9

---

**QUESTION 8**

What is the minimum required number of hosts to provide data redundancy for a vSAN stretched cluster using dual-site mirroring and local protection with 1 failure - RAID-1 (Mirroring)?

- A. 3 hosts
- B. 3 hosts
- C. 4 hosts
- D. 6 hosts

Correct Answer: D

Explanation: The minimum required number of hosts to provide data redundancy for a vSAN stretched cluster using dual-site mirroring and local protection with 1 failure - RAID-1 (Mirroring) is six hosts. This is because a vSAN stretched cluster requires at least three hosts per site, and each site must have enough hosts to tolerate one host failure. Therefore, the minimum configuration is three hosts per site, plus one witness host at a third site, for a total of six hosts. References: [VMware vSAN Specialist v2 EXAM 5V0- 22.23], page 14

---

**QUESTION 9**

A vSAN administrator is planning to deploy a new vSAN cluster with these requirements:

Physical adapters share capacity among several traffic types  
Guaranteed bandwidth for vSAN during bandwidth contention  
Enhanced security



Which two actions should be taken to configure the new vSAN cluster to meet these requirements? (Choose two.)

- A. Create static routes between the vSAN hosts
- B. Use IOPS Limit rules in storage policies
- C. Utilize Network I/O Control
- D. Enable jumbo frames
- E. Isolate vSAN traffic in a VLAN

Correct Answer: CE

Explanation: Utilizing Network I/O Control and isolating vSAN traffic in a VLAN are the two actions that should be taken to configure the new vSAN cluster to meet the requirements. Network I/O Control allows the vSAN administrator to create network resource pools and assign bandwidth shares or reservations to different traffic types, such as vSAN, vMotion, or management. This ensures that vSAN traffic has guaranteed bandwidth during contention and can achieve better performance and availability. Isolating vSAN traffic in a VLAN enhances the security of the cluster by preventing unauthorized access or interference from other network segments. It also simplifies the network configuration and management by reducing the broadcast domain and avoiding IP address conflicts. Creating static routes between the vSAN hosts, using IOPS Limit rules in storage policies, and enabling jumbo frames are not necessary or recommended actions for this scenario. Static routes are not required for vSAN communication, as vSAN uses multicast or unicast depending on the version and configuration. IOPS Limit rules are used to limit the IOPS allocated to an object, which can degrade the performance and latency of the application. Jumbo frames can improve the network efficiency and throughput, but they are not mandatory for vSAN and require consistent configuration across all network devices.

References: Network I/O Control vSAN Network Design Guide

## QUESTION 10

A site administrator wishes to implement HCI mesh between two clusters on vSAN that are located in geographically separate sites and which are administered within a single datacenter.

Which two requirements should the vSAN administrator consider to accomplish this goal? (Choose two.)

- A. Either Layer 2 or Layer 3 communications can be used
- B. A leaf spine topology is required for core redundancy and reduced latency
- C. NIC teaming must be implemented for the vSAN network vmkernel port
- D. The configuration must meet the same latency and bandwidth requirement as local vSAN
- E. Encryption must be disabled prior to configuring HCI mesh

Correct Answer: AD

Explanation: To implement HCI mesh between two clusters on vSAN that are located in geographically separate sites, the vSAN administrator should consider the following requirements: Either Layer 2 or Layer 3 communications can be used. HCI mesh supports both Layer 2 and Layer 3 network configurations, as long as the network latency and bandwidth requirements are met. The configuration must meet the same latency and bandwidth requirement as local vSAN. HCI mesh requires a network latency of less than or equal to 5 ms RTT between any two hosts in the participating clusters, and a network bandwidth of at least 10 Gbps for the vSAN network vmkernel port. References: 3: VMware vSAN Specialist v2 Exam Preparation Guide, page 15

**QUESTION 11**

An all-flash vSAN ESA cluster contains four nodes.

Which two storage policies can the cluster satisfy? (Choose two.)

- A. FTT=3 (RAID-1 Mirroring)
- B. FTT=2 (RAID-1 Mirroring)
- C. FTT=1 (RAID-5 Erasure Coding)
- D. FTT=1 (RAID-1 Mirroring)
- E. FTT=2 (RAID-6 Erasure Coding)

Correct Answer: BE

Explanation: An all-flash vSAN ESA cluster with four nodes can satisfy the storage policies that require FTT=2 (RAID-1 Mirroring) or FTT=2 (RAID-6 Erasure Coding). These policies mean that the cluster can tolerate two host failures while maintaining data availability and redundancy. RAID-1 Mirroring creates three replicas of each object across different hosts, while RAID-6 Erasure Coding splits each object into four data segments and two parity segments across different hosts. Both policies require at least four hosts in the cluster to meet the FTT=2 requirement. The other options are not correct. An all-flash vSAN ESA cluster with four nodes cannot satisfy the storage policies that require FTT=3 (RAID-1 Mirroring) or FTT=1 (RAID-5 Erasure Coding). These policies mean that the cluster can tolerate three or one host failure respectively, but they require more or less hosts than four to do so. RAID-1 Mirroring with FTT=3 requires at least six hosts in the cluster to create four replicas of each object, while RAID-5 Erasure Coding with FTT=1 requires at least three hosts in the cluster to split each object into two data segments and one parity segment. References: vSAN Express Storage Architecture; RAID Configurations, FTT, and Host Requirements

---

**QUESTION 12**

The Resyncing Objects view in the vCenter UI reports that some objects are currently resyncing. Which two actions would cause this situation? (Choose two.)

- A. A change to the storage policy is applied to the objects.
- B. DRS is relocating VMs between vSAN nodes.
- C. A host failure occurs in the cluster
- D. HA Virtual Machine Monitoring forced a VM to reboot.
- E. VM snapshot is being deleted.

Correct Answer: AC

Explanation: Two actions that would cause some objects to be currently resyncing are: A change to the storage policy is applied to the objects. This action triggers a resynchronization of objects to make them compliant with the new policy settings, such as FTT, RAID level, stripe width, etc. The resynchronization process copies data from one host to another to create or update replicas or parity segments. A host failure occurs in the cluster. This action causes some objects to become non-compliant with their storage policy, as they lose one or more replicas or parity segments due to the host failure. The resynchronization process rebuilds the missing components on other hosts in the cluster to restore





compliance and availability. References: : VMware vSphere Storage Guide, page 129 : Monitor the Resynchronization Tasks in the vSAN Cluster 1 : VMware vSAN Specialist v2 Exam Preparation Guide, page 13

---

### QUESTION 13

An administrator has 24 physical servers that need to be configured with vSAN. The administrator needs to ensure that a single rack failure is not going to affect the data availability. The number of racks used should be minimized.

What has to be done and configured to achieve this goal?

- A. Distribute servers across at least two different racks and configure two fault domains
- B. Configure disk groups with a minimum of four capacity disks in each server and distribute them across four racks
- C. Enable deduplication and compression
- D. Distribute servers across at least three different racks and configure three fault domains

Correct Answer: D

Explanation: To ensure that a single rack failure is not going to affect the data availability, while minimizing the number of racks used, the administrator has to do the following: Distribute servers across at least three different racks. This is because vSAN supports up to three fault domains per cluster, which can be used to tolerate one or two failures. If only two racks are used, then only one failure can be tolerated. Configure three fault domains. A fault domain is a logical grouping of hosts that share a common failure point, such as a rack or a power supply. By configuring fault domains, vSAN can place replicas of an object across different fault domains, so that a failure within one fault domain does not result in data loss or unavailability. References: 4: VMware vSAN Specialist v2 Exam Preparation Guide, page 13

---

### QUESTION 14

During yesterday's business hours, a cache drive failed on one of the vSAN OSA nodes. The administrator reached out to the manufacturer and received a replacement drive the following day. When the drive failed, vSAN started a resync to ensure the health of the data, and all objects are showing a healthy and compliant state. The vSAN administrator needs to replace the failed cache drive.

Which set of steps should the vSAN administrator take?

- A. Physically replace the failed cache device, and vSAN will automatically create a new disk group. Then, remove the disk group with the failed device.
- B. Place the disk group into maintenance mode, and select Full Data Migration. Then, physically replace the failed cache device. Afterwards, vSAN will rebuild the disk group automatically.
- C. Remove the existing vSAN disk group and physically replace the device. Then check to verify that the ESXi host automatically detects the new device. Afterwards manually recreate the Disk Group.
- D. Physically replace the failed cache device, and vSAN will automatically allocate the storage. Then, rebalance the cache layer.

Correct Answer: C

Explanation: To replace a failed cache drive in a vSAN OSA cluster, the vSAN administrator should remove the existing vSAN disk group and physically replace the device. Then check to verify that the ESXi host automatically detects the





new device Afterwards manually recreate the Disk Group. This is because when a cache drive fails, it affects the entire disk group that contains it, and vSAN does not allow removing only the cache drive from a disk group. Therefore, the administrator must remove the whole disk group before replacing the cache drive, and then recreate it with the new cache drive and the existing capacity drives. The other options are not correct. Physically replacing the failed cache drive without removing the disk group first might cause errors or inconsistencies in vSAN configuration. vSAN will not automatically create a new disk group or allocate storage after replacing a cache drive, as these actions require manual intervention from the administrator. Rebalancing the cache layer is not necessary after replacing a cache drive, as vSAN will automatically distribute data across all devices in the disk group. References: Replace a Flash Caching Device on a Host; How to manually remove and recreate a vSAN disk group using esxcli

### QUESTION 15

Refer to the exhibit.

An administrator uses SSH to log into a vSAN ESA host and runs the `esxcli vsan debug object overview` command.

Object UUID	Group UUID	Version	Size	Used	SPBM Profile	Healthy Components
49413f63-84bd-4aba-2ba6-0050560659c0	1a413f63-a8d1-fa7b-0809-0050560659c0	17	0.12 GB	0.01 GB	vSAN Default Storage Policy	7 of 8
1c413f63-4c1a-73bc-904e-0050560659c0	1a413f63-a8d1-fa7b-0809-0050560659c0	17	2.00 GB	0.70 GB	vSAN Default Storage Policy	5 of 8
1a413f63-a8d1-fa7b-0809-0050560659c0	1a413f63-a8d1-fa7b-0809-0050560659c0	17	255.00 GB	0.05 GB	vSAN Default Storage Policy	4 of 8
cf403f63-ec5-da41-8599-005056065997	cf403f63-ec5-da41-8599-005056065997	17	255.00 GB	0.04 GB	vSAN Default Storage Policy	7 of 8
d0403f63-f7af-45cd-2e8a-005056065997	cf403f63-ec5-da41-8599-005056065997	17	2.00 GB	0.70 GB	vSAN Default Storage Policy	7 of 8
ef403f63-fe7b-6620-9d4f-005056065997	cf403f63-ec5-da41-8599-005056065997	17	0.12 GB	0.01 GB	vSAN Default Storage Policy	5 of 8
db413f63-4ca4-7882-1b50-005056065979	db413f63-4ca4-7882-1b50-005056065979	17	255.00 GB	0.12 GB	vSAN ESA Default Policy - RAID5	8 of 8
dd413f63-e0e3-929d-9b93-005056065979	db413f63-4ca4-7882-1b50-005056065979	17	90.00 GB	0.01 GB	vSAN ESA Default Policy - RAID5	5 of 8
e2413f63-4072-62cf-2077-005056065979	db413f63-4ca4-7882-1b50-005056065979	17	4.00 GB	0.01 GB	vSAN ESA Default Policy - RAID5	8 of 8
f0403f63-e677-850f-db46-005056065979	f0403f63-e677-850f-db46-005056065979	17	255.00 GB	0.05 GB	vSAN Default Storage Policy	7 of 8
0d413f63-8c52-b213-3866-005056065979	f0403f63-e677-850f-db46-005056065979	17	0.12 GB	0.01 GB	vSAN Default Storage Policy	4 of 8
f1403f63-365f-559e-8165-005056065979	f0403f63-e677-850f-db46-005056065979	17	2.00 GB	0.72 GB	vSAN Default Storage Policy	4 of 8
f4403f63-50e3-85c4-ed42-0050560659b4	f4403f63-50e3-85c4-ed42-0050560659b4	17	255.00 GB	3.54 GB	vSAN Default Storage Policy	5 of 9

The administrator notices the Healthy Components column, the last column, is reporting some components are not in a fully healthy state.

What could cause this behavior?

- A. New physical disks have been claimed and a rebalance operation is underway.
- B. The applied Storage policy has been updated.
- C. New VMDKs have been added to multiple VMs, but the storage policy has not finished applying.
- D. One host is in maintenance mode with ensure accessibility.

Correct Answer: D

Explanation: The most likely cause for some components to be not in a fully healthy state is that one host is in maintenance mode with the ensure accessibility option. This option creates temporary durability components on other hosts to maintain the required number of failures to tolerate (FTT) until the original components are restored or rebuilt. These durability components are not considered fully healthy because they do not have full redundancy and might not be compliant with the storage policy. The other options do not explain why some components are not fully healthy, as they do not affect the FTT or the compliance state of the objects. References: Durability Components; esxcli vsan debug object overview