

# 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.certbus.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

A three-node vSAN OSA cluster with business critical intensive I/O workload is running out of capacity. Each host consists of five disk groups with four capacity disks. The administrator needs to expand the capacity of the vSAN datastore as soon as possible.

What should the administrator do?

- A. Enable Deduplication and Compression on the cluster level
- B. Add additional capacity by adding a disk on one host and creating a storage pool
- C. Add additional capacity by adding a vSAN ReadyNode to the cluster
- D. Add additional capacity disks to each disk group

Correct Answer: D

Explanation: The correct answer is D, add additional capacity disks to each disk group. This is because adding capacity disks to existing disk groups is the fastest and easiest way to expand the capacity of the vSAN datastore without disrupting any ongoing operations or requiring additional hardware. The administrator can add up to five capacity disks per disk group in vSAN OSA, which means each host can have up to 25 capacity disks in total. The administrator should make sure that the new capacity disks are unformatted and not partitioned, so that vSAN can recognize and claim them. The administrator should also manually rebalance the cluster after adding the capacity disks to distribute the data evenly across the new devices. The other options are incorrect for the following reasons: A, enable Deduplication and Compression on the cluster level, is incorrect because enabling Deduplication and Compression is not a recommended way to expand the capacity of the vSAN datastore. Deduplication and Compression is a space efficiency feature that reduces the logical space consumption of data by eliminating duplicate blocks and applying compression algorithms. However, enabling Deduplication and Compression requires a full data evacuation and resynchronization, which can be disruptive and time-consuming. Deduplication and Compression also introduces additional CPU and memory overhead, which can affect the performance of the cluster. Deduplication and Compression is only supported on all-flash clusters, not on hybrid clusters. B, add additional capacity by adding a disk on one host and creating a storage pool, is incorrect because creating a storage pool is not supported in vSAN OSA. A storage pool is a new configuration introduced in vSAN 8 ESA, where all disks are treated as capacity disks and use a new algorithm to distribute data across them. This configuration is not compatible with vSAN OSA, which uses a disk group configuration where one disk is designated as a cache disk and the rest are capacity disks. To use a storage pool, the administrator would need to migrate to vSAN 8 ESA on a new cluster with new hardware. C, add additional capacity by adding a vSAN ReadyNode to the cluster, is incorrect because adding a vSAN ReadyNode to the cluster is not the fastest or easiest way to expand the capacity of the vSAN datastore. A vSAN ReadyNode is a preconfigured server that meets the hardware requirements for running vSAN. Adding a vSAN ReadyNode to the cluster would require additional hardware procurement, installation, and configuration. It would also increase the compute capacity of the cluster, which may not be necessary for the workload. Adding a vSAN ReadyNode would also trigger a resynchronization of data across the cluster, which can affect the performance and availability of the cluster. References: VMware vSAN Specialist v2 Exam Preparation Guide, page 10

## QUESTION 2

In which type of environment is vSAN storage used as a mandatory, primary storage?

- A. VMware Cloud on AWS
- B. VMware Horizon
- C. VMware Aria Automation

#### D. TanzuKubernetes Grid Integrated Edition

Correct Answer: A

Explanation: VMware Cloud on AWS is a service that delivers a fully managed VMware SDDC on AWS infrastructure. It uses vSAN as the mandatory, primary storage for the SDDC clusters. vSAN provides a high-performance, resilient, and secure shared storage solution for the VMware Cloud on AWS environment. The other options are not correct, as vSAN is not mandatory or primary for them. VMware Horizon, VMware Aria Automation, and Tanzu Kubernetes Grid Integrated Edition can use vSAN as an optional or secondary storage solution, but they can also use other types of storage. References: , section 1.1; , section 1.2

### QUESTION 3

A vSAN administrator was presented with 30 additional vSAN ReadyNodes to add to an existing vSAN cluster. There is only one administrator to complete this task.

What is the fastest approach?

- A. Run vim-cmd to capture, and apply the configuration from an existing host
- B. Launch Quickstart to Add Hosts to a vSAN Cluster
- C. Clone the ESXi boot partition to all new hosts, since the hardware is identical
- D. Use a Host Profile that was extracted from an existing host

Correct Answer: D

Explanation: To add 30 additional vSAN ReadyNodes to an existing vSAN cluster with the fastest approach, the vSAN administrator should use a Host Profile that was extracted from an existing host. A Host Profile is a configuration template that captures the settings of a reference host and applies them to other hosts or clusters. This way, the administrator can quickly and consistently configure multiple hosts with the same settings, such as network, storage, security, and services. The other options are not correct. Running vim-cmd to capture and apply the configuration from an existing host is not as fast or convenient as using a Host Profile, as it requires running commands on each host individually. Launching Quickstart to Add Hosts to a vSAN Cluster is not possible, as Quickstart is only available for new clusters or clusters that were configured through Quickstart. Cloning the ESXi boot partition to all new hosts is not recommended, as it might cause conflicts or errors with the host identity, network settings, or licenses. References: Configuring Hosts Using Host Profile; Using Quickstart to Configure and Expand a vSAN Cluster

### QUESTION 4

An administrator wants to assign a storage policy to a workload on a two-node vSAN OSA cluster consisting of three disk groups each with nested fault domains. The virtual machine must be protected against a disk or disk group failure.

Which two storage policies meet these requirements? (Choose two.)

- A. RAID-5/FTT 2
- B. RAID-1/FTT 3
- C. RAID-6/FTT 2
- D. RAID-5/FTT 1

E. RAID-1/FTT 1

Correct Answer: CE

Explanation: To protect a virtual machine against a disk or disk group failure, the storage policy must have a failure tolerance method (FTM) of RAID-1 or RAID-6 and a failure to tolerate (FTT) value of at least 1. RAID-1 mirrors the data across multiple disk groups, while RAID-6 uses erasure coding to stripe the data and parity information across multiple disk groups. RAID-5 is not suitable for this scenario, as it can only tolerate one disk failure per stripe. FTT 2 or 3 would require more disk groups than available in the cluster. Therefore, the correct options are C and E. References: 1, page 8; 2, section 3.1

## QUESTION 5

An organization wants to implement a virtual desktop infrastructure (VDI) solution on their vSAN storage. They also need to store their applications running inside the VDI environment on vSAN storage. Which two end-user computing (EUC) solutions could be implemented to satisfy the requirements of the organization? (Choose two.)

- A. Agp\_ Volumes
- B. Workspace ONE Access
- C. Horizon
- D. Workspace ONE UEM
- E. Dynamic Environment Manager

Correct Answer: CE

Explanation: Horizon and Dynamic Environment Manager are two end-user computing (EUC) solutions that can be implemented on vSAN storage to provide a virtual desktop infrastructure (VDI) solution and store applications running inside the VDI environment. Horizon is a platform that delivers virtual desktops and applications across a variety of devices and locations, while Dynamic Environment Manager is a tool that provides personalization and dynamic policy configuration across any virtual, physical, and cloud- based Windows desktop environment. The other solutions are not directly related to VDI or application storage on vSAN. References: VMware vSAN Specialist v2 EXAM 5V0-22.23, page 8, Objective 3.5; [Horizon]; [Dynamic Environment Manager]

## QUESTION 6

A vSAN administrator has two identical VMware vSAN clusters, one for staging workloads and another for production workloads. Due to an unforeseen capacity requirement, the vSAN administrator is tasked with merging the staging vSAN cluster into the production.

Which three actions should the vSAN administrator perform on the staging cluster prior to moving the vSAN nodes to the production cluster? (Choose three.)

- A. Disable vSAN Services
- B. Delete all Disk Groups
- C. Enable File Services
- D. Delete all partitions from the capacity disks

E. Mark the disks for partial reservation

F. Remove all capacity drives

Correct Answer: ABD

Explanation: The three actions that the vSAN administrator should perform on the staging cluster prior to moving the vSAN nodes to the production cluster are: Disable vSAN Services: This will stop any vSAN-related operations on the staging cluster, such as resynchronization, rebalancing, or repair. This will also prevent any new virtual machines from being created or migrated to the staging cluster. Delete all Disk Groups: This will remove all disks from the vSAN cluster and erase all data on them. This will also free up the disks for use in the production cluster. Delete all partitions from the capacity disks: This will ensure that there are no remnants of any previous vSAN configuration on the disks. This will also avoid any potential conflicts or errors when adding the disks to the production cluster. Enabling File Services, marking the disks for partial reservation, and removing all capacity drives are not necessary or recommended actions for this scenario. Enabling File Services would add an unnecessary layer of complexity and overhead to the staging cluster. Marking the disks for partial reservation would reduce the available capacity for vSAN and potentially cause performance issues. Removing all capacity drives would leave only cache disks in the staging cluster, which would not be compatible with vSAN. References: VMware vSAN Specialist v2 Exam Preparation Guide, page 10

## QUESTION 7

A customer wishes to host a new range of applications with high-performance needs, specifically, low latency.

The applications are required to be hosted at company-owned edge locations, each with minimal rack space (three host slots per edge location for this project).

Which deployment options would satisfy the customer's needs, while maximizing the amount of capacity available per deployment?

A. A new three-node vSAN 8.0 All-Flash Cluster with OSA in each edge location Each application VM configured with a RAID-5 VM storage policy

B. A new three-node vSAN 8.0 All-Flash Cluster with OSA in each edge location Each application VM configured with a RAID-1 VM storage policy

C. A new three-node vSAN 8.0 All-Flash Cluster with ESA in each edge location Each application VM configured with a RAID-1 VM storage policy

D. A new three-node vSAN 8.0 All-Flash Cluster with ESA in each edge location Each application VM configured with a RAID-5 VM storage policy

Correct Answer: B

Explanation: To satisfy the customer's needs for high-performance, low-latency applications at edge locations, the best deployment option is to use a new three-node vSAN 8.0 All-Flash Cluster with OSA in each edge location and configure each application VM with a RAID-1 VM storage policy. This option will provide the following benefits: All-flash clusters offer the highest performance and lowest latency for vSAN, as they use flash devices for both cache and capacity tiers. Flash devices have faster read and write operations than magnetic disks, and they also support advanced features such as deduplication, compression, and encryption. OSA stands for One Socket Architecture, which means that each host has only one CPU socket with multiple cores. This reduces the licensing cost and complexity of vSphere and vSAN, as well as the power consumption and cooling requirements of the hosts. OSA also improves the performance of vSAN by eliminating the NUMA effect, which is the latency caused by accessing memory or devices across different CPU sockets. RAID-1 is a mirroring technique that creates two copies of each data component and places them on different hosts. This provides high availability and fault tolerance for the application VMs, as they can survive the failure of one host or disk. RAID-1 also offers better performance than RAID-5 or RAID-6, as it does not incur any parity overhead or

additional write operations. The other options are not optimal for the customer's needs, as they either sacrifice performance or capacity. Option A uses RAID-5, which is an erasure coding technique that splits each data component into three data segments and one parity segment, and distributes them across four hosts. This reduces the capacity consumption by 25%, but it also increases the write latency and network traffic, as each write operation requires four hosts to participate. Option C uses ESA, which stands for Enterprise Storage Architecture, which means that each host has two CPU sockets with multiple cores. This increases the licensing cost and complexity of vSphere and vSAN, as well as the power consumption and cooling requirements of the hosts. ESA also introduces the NUMA effect, which can degrade the performance of vSAN by adding latency to access memory or devices across different CPU sockets. Option D uses RAID-5 with ESA, which combines the disadvantages of both options A and C.

## QUESTION 8

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 9

A vSAN administrator has a vSAN cluster that is using vSphere Lifecycle Manager (vLCM) to manage hypervisor, server drivers, and firmware. All hosts in the cluster are compliant according to the vLCM image.

A 10GB NIC on the servers is experiencing issues, and the vSAN administrator determines a new network driver will resolve the problem. Unfortunately, the required NIC driver is a newer version compared to the driver provided by the most recent Vendor Add-on.

Which action should the vSAN administrator take to ensure the latest network driver is installed on the NIC before remediation?

- A. Make sure the vLCM image is configured to use the most recent version of the Vendor Add-on
- B. Add an individual component to the vLCM image that has the updated NIC driver
- C. Remove the Vendor Add-on from the vLCM image, and then manually install the network driver on the servers

D. Modify the vLCM image to omit the NIC Driver, and then manually update the servers with the required NIC driver

Correct Answer: B

Explanation: To ensure the latest network driver is installed on the NIC before remediation, the vSAN administrator should add an individual component to the vLCM image that has the updated NIC driver. This action allows the administrator to override the driver provided by the vendor add-on and use a newer version that is compatible with the ESXi version and the hardware device. The administrator can add an individual component to the vLCM image by importing it from a ZIP file or selecting it from the vLCM depot. The other options are not correct. Making sure the vLCM image is configured to use the most recent version of the vendor add-on will not help, as the required NIC driver is a newer version than the one provided by the vendor add-on. Removing the vendor add-on from the vLCM image or modifying the vLCM image to omit the NIC driver will not ensure the latest network driver is installed on the NIC, as these actions will leave the NIC without any driver update. Manually installing or updating the network driver on the servers is not recommended, as it might cause inconsistency and non-compliance in the vLCM image. References: vSphere Lifecycle Manager Image Components; [Add an Individual Component to an Image]

## QUESTION 10

A customer wishes to host a new range of applications with high-performance requirements, specifically, low latency. The current vSAN platform is based on ReadyNode hardware and uses a vSAN 7.0 U2 hybrid topology configuration.

Which would satisfy the customer's requirement?

- A. Deploy the application on a new cluster with vSAN 8.0 ESA using a new hardware design
- B. Deploy the new applications on the existing cluster with a RAID-6 VM storage policy and an additional stripe width of 4
- C. Deploy the application on a new cluster with vSAN 8.0 OSA using the existing hybrid configuration
- D. Perform an in-place upgrade from vSAN 7.0 U2 OSA to vSAN 8.0 ESA

Correct Answer: A

Explanation: Deploying the application on a new cluster with vSAN 8.0 ESA using a new hardware design is the correct answer because it will satisfy the customer's requirement for low latency. vSAN 8.0 ESA is a new architecture that uses a storage pool configuration where all disks are treated as capacity disks and use a new algorithm to distribute data across them. This improves the I/O flow, reduces the write amplification, and eliminates the cache tier bottleneck. Using a new hardware design with all-flash disks or NVMe disks will further enhance the performance and latency of the application, as these disks have faster read and write speeds than hybrid disks. Deploying the new applications on the existing cluster with a RAID-6 VM storage policy and an additional stripe width of 4, deploying the application on a new cluster with vSAN 8.0 OSA using the existing hybrid configuration, and performing an in-place upgrade from vSAN 7.0 U2 OSA to vSAN

8.0 ESA are not valid or optimal solutions for this scenario. Deploying the new applications on the existing cluster with a RAID-6 VM storage policy and an additional stripe width of 4 will increase the resiliency and availability of the data, but it will also increase the network traffic, disk space consumption, and parity calculation overhead, which will negatively affect the latency and performance of the application. Deploying the application on a new cluster with vSAN 8.0 OSA using the existing hybrid configuration will not improve the latency significantly, as vSAN 8.0 OSA still uses the same disk group configuration as vSAN 7.0 U2 OSA, where one disk is designated as a cache disk and the rest are capacity disks. The cache disk can still become a bottleneck for high-performance applications, especially if it is not an SSD or NVMe disk. Performing an in-place upgrade from vSAN 7.0 U2 OSA to vSAN 8.0 ESA is not possible, as vSAN ESA requires a different hardware design than vSAN OSA. The existing disk groups need to be deleted and all disks need to be erased before switching to vSAN ESA. References: [VMware vSAN Specialist v2 Exam Preparation Guide], page 6 What's New in VMware vSAN 8.0

### QUESTION 11

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 12

An application refactor requires significant storage that is being added for logs stored on a VM vDISK. The application VMs run on a dedicated vSAN enabled vSphere Cluster with custom CPUs and RAM, and therefore, cannot vMotion to another vSAN enabled cluster.

The administrator needs a vSAN feature that can be used to allocate additional storage from another vSAN enabled vSphere cluster to this vSAN enabled Cluster.

Which vSAN feature should be used for this purpose?

- A. vSAN File Services
- B. vSAN HCI Mesh
- C. vSAN Replication
- D. vSAN Stretched Clusters

Correct Answer: B

Explanation: To allocate additional storage from another vSAN enabled vSphere cluster to this vSAN enabled Cluster, the administrator should use the vSAN HCI Mesh feature. This feature allows a vSAN cluster to consume storage resources from another vSAN cluster without requiring the hosts to be part of the same cluster. This way, the administrator can leverage the unused or underutilized storage capacity from another cluster and avoid purchasing new

hardware or migrating VMs. The vSAN HCI Mesh feature also supports storage policies, encryption, deduplication and compression, and erasure coding across clusters

12 References: 1: VMware vSAN Specialist v2 Exam Preparation Guide, page 15

2: VMware vSAN 7 Update 1 - HCI Mesh 3

### QUESTION 13

An administrator is troubleshooting a vSAN performance issue. In the vSAN performance monitor there is a high latency on the vSAN cluster.

What is a possible cause of this?

- A. The Virtual Machines are using PVSCSI controllers.
- B. Erasure Coding is disabled in the storage policy.
- C. There is congestion in one or more disk groups.
- D. Jumbo frames are not enabled on the VMkernel adapters.

Correct Answer: C

Explanation: A possible cause of high latency on the vSAN cluster is that there is congestion in one or more disk groups. Congestion is a measure of how busy the storage devices are in handling I/O requests. When congestion is high, it means that the storage devices are overloaded and cannot process the requests fast enough, resulting in increased latency and reduced throughput. Congestion can be caused by various factors, such as insufficient cache capacity, disk failures, network issues, or heavy workload. The other options are not likely to cause high latency on the vSAN cluster. The Virtual Machines can use PVSCSI controllers without affecting latency, as they are optimized for high performance. Erasure Coding is a space efficiency feature that does not impact latency significantly. Jumbo frames are not required for vSAN, and enabling them does not guarantee lower latency. References: vSAN Performance Monitor; [vSAN Congestion Explained]

### QUESTION 14

A customer has deployed a new vSAN cluster with the following configuration:

5 x vSAN ReadyNodes

All Flash

12 TB Raw Storage

vSAN 8 is deployed with ESA.

New VMs are configured with a RAID-5 VM policy.

Which statement is accurate?

- A. vSAN will use a 2+1 RAID-5 data placement scheme with parity will be used
- B. RAID 5 will provide an FTT=2 level of protection in this case

C. vSAN will use a 4+1 RAID-5 data placement scheme with parity will be used

D. vSAN will spread the components across all of the disk groups

Correct Answer: C

Explanation: vSAN will use a 4+1 RAID-5 data placement scheme with parity will be used is the correct answer because vSAN 8 ESA uses adaptive RAID-5 erasure coding that depends on the number of hosts in the cluster. If the cluster has 6 or more hosts, vSAN will use a 4+1 RAID-5 scheme, where the data is written as a stripe of 4 data bits and 1 parity bit across 5 hosts. This provides a failure tolerance of 1 (FTT=1) and a space efficiency of 1.25x. If the cluster has less than 6 hosts (3 to 5), vSAN will use a 2+1 RAID-5 scheme, where the data is written as a stripe of 2 data bits and 1 parity bit across 3 hosts. This also provides a failure tolerance of 1 (FTT=1) but a space efficiency of 1.5x. In this case, the cluster has 5 hosts, so vSAN will use the 4+1 RAID-5 scheme. The other options are incorrect for the following reasons: A, vSAN will use a 2+1 RAID-5 data placement scheme with parity will be used, is incorrect because vSAN will only use this scheme if the cluster has less than 6 hosts but more than 2 hosts. In this case, the cluster has 5 hosts, so vSAN will use the 4+1 RAID-5 scheme. B, RAID 5 will provide an FTT=2 level of protection in this case, is incorrect because RAID 5 can only provide an FTT=1 level of protection, regardless of the number of hosts or the data placement scheme. To achieve an FTT=2 level of protection, vSAN would need to use RAID 6 erasure coding, which requires at least 6 hosts in the cluster. D, vSAN will spread the components across all of the disk groups, is incorrect because vSAN will not necessarily spread the components across all of the disk groups in the cluster. vSAN will only spread the components across as many disk groups as needed to meet the storage policy requirements and to balance the load and capacity. In this case, vSAN will only need to spread the components across 5 disk groups for each stripe of RAID-5 data. References: VMware vSAN Specialist v2 Exam Preparation Guide, page 11 Adaptive RAID-5 Erasure Coding with the Express Storage Architecture in vSAN 8

## QUESTION 15

A vSAN administrator needs to update vSAN from version 7.0.2 to version 8.0. Which is the correct order to perform the update?

A. vSphere -> vCenter -> vSAN on-disk format

B. vSphere -> vSAN on-disk format -> vCenter

C. vCenter -> vSphere -> vSAN on-disk format

D. vSAN on-disk format -> vSphere -> vCenter

Correct Answer: C

Explanation: The correct order to perform the update from vSAN version 7.0.2 to version 8.0 is to upgrade the vCenter Server first, then upgrade the ESXi hosts, and finally upgrade the vSAN on-disk format. This order follows the general vSphere upgrade order, which ensures compatibility and interoperability between different components. Upgrading the vCenter Server first allows it to manage and monitor the ESXi hosts and the vSAN cluster during the upgrade process. Upgrading the ESXi hosts second ensures that they have the latest software patches and drivers for vSAN. Upgrading the vSAN on-disk format last enables the new features and functionality of vSAN 8.0. The other options are not correct, as they do not follow the recommended upgrade order.

[Latest 5V0-22.23 Dumps](#)

[5V0-22.23 VCE Dumps](#)

[5V0-22.23 Braindumps](#)