



5V0-22.23^{Q&As}

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.geekcert.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 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 2

A vSAN administrator is using the vSAN ReadyNode Sizer to build a new environment. While entering the cluster configurations, a fellow colleague inquires about the Operations Reserve option.

What is the purpose of using this option?

- A. Provides space for internal operations
- B. Configures space for external operations
- C. Reserves space for tolerating failures
- D. Allocates space for vSAN upgrades

Correct Answer: A

Explanation: The purpose of using the Operations Reserve option in the vSAN ReadyNode Sizer is to provide space for internal operations such as deduplication, compression, encryption, snapshots, clones, and rebalancing. The Operations Reserve is calculated as a percentage of the total usable capacity of the vSAN cluster. The default value is 30%, but it can be adjusted based on the expected workload characteristics and data services requirements. The other options are not correct, as they do not describe the Operations Reserve option. Configuring space for external operations, reserving space for tolerating failures, and allocating space for vSAN upgrades are not part of the



Operations Reserve option. References: 2, section 2; , section 3

QUESTION 3

A customer has deployed a new vSAN Cluster with the following configuration: 6 x vSAN ReadyNodes All Flash 12 TB Raw Storage

vSAN 8 is deployed with ESA.

VMs are configured with a RAID-5 VM policy.

During failure testing, before the new platform is placed into production one of the ESXi hosts is made unavailable.

Which RAID-5 data placement schemes will vSAN use with this failure condition?

- A. vSAN can protect the platform using adaptive RAID 5 if the ESXi host fails to return
- B. VMware HA will migrate the storage objects to another node in the cluster
- C. Some VM data will be unavailable until the failed ESXi host is recovered
- D. The data components on the hosts will be marked as degraded

Correct Answer: D

Explanation: When a host in a vSAN stretched cluster goes offline, the data components on the hosts will be marked as degraded. This means that the data is still available, but the redundancy level is reduced. vSAN will try to rebuild the missing components on another host in the same fault domain, if there is enough capacity and resources. If the host comes back online within 60 minutes, vSAN will resync the data and restore the redundancy level. If the host does not come back online within 60 minutes, vSAN will rebuild the missing components on another fault domain, if there is enough capacity and resources. This will incur additional network traffic across the witness link. References: VMware vSAN Specialist v2 EXAM 5V0-22.23, page 17

QUESTION 4

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 5

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

[5V0-22.23 PDF Dumps](#)[5V0-22.23 VCE Dumps](#)[5V0-22.23 Study Guide](#)