

# VMware

## Exam Questions 5V0-22.23

VMware vSAN Specialist (v2)



#### NEW QUESTION 1

A vSAN administrator has a cluster configured with a Storage Pool that was moved to a new physical DC. Upon checking on the vSAN cluster health status, one of the ESXi hosts has two storage devices in a degraded state and must be replaced. What must the vSAN administrator do to restore the health of the vSAN cluster with minimum risk?

- A. Remove the host from ySAN configuration, replace the faulty disks, re-create the storage pool
- B. Remove the entire storage pool, install the new devices, re-create the storage pool
- C. Remove the host from the cluster, replace the faulty disks, re-add the host to the cluster
- D. Remove the devices from the storage pool, replace the storage devices, claim the new devices in vSAN

**Answer: D**

#### Explanation:

To restore the health of the vSAN cluster with minimum risk, the vSAN administrator must remove the devices from the storage pool, replace the storage devices, and claim the new devices in vSAN. This is because removing and replacing devices in a storage pool does not affect the availability or performance of the objects stored in that pool. The storage pool automatically rebalances the objects across the remaining devices in the pool when a device is removed, and distributes the objects across the new devices when they are added. This process is faster and safer than removing and re-adding a host to the cluster, which requires resynchronization of all objects on that host<sup>4</sup> References: 4: VMware vSphere Storage Guide, page 133 : VMware vSAN Design and Sizing Guide, page 38

#### NEW QUESTION 2

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.

**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

#### NEW QUESTION 3

A host in a vSAN stretched cluster goes offline during an unplanned event. Which action will be triggered from AQC on the vSAN cluster?

- A. AQC will trigger a vMotion of VMs that went offline.
- B. AQC will restart the VMs that went offline.
- C. AQC will recalculate the quorum on an object.
- D. AQC will create a vSAN alarm.

**Answer: C**

#### Explanation:

When a host in a vSAN stretched cluster goes offline, vSAN will use Adaptive Resync to recalculate the quorum on an object. Quorum is the minimum number of votes that an object needs to be available. For example, a RAID-1 object with two data components and one witness component needs two votes out of three to be available. If one data component goes offline, the object still has quorum and is available. However, if both data components go offline, the object loses quorum and is unavailable. Adaptive Resync will adjust the quorum requirement based on the availability of components and fault domains. For example, if one fault domain goes offline, Adaptive Resync will lower the quorum requirement to one vote out of two, so that the object can remain available with one data component and one witness component. References: VMware vSAN Specialist v2 EXAM 5V0-22.23, page 18

#### NEW QUESTION 4

An administrator has to perform maintenance on one of the hosts in a three-node vSAN Cluster. Which maintenance mode option will give the administrator the best availability for the VMs with the least effort and data transfer?

- A. Migrate all VMs and their storage from the host to a different storage system
- B. Full data migration
- C. Migrate all VMs and their storage from the host to a different vSphere cluster
- D. Ensure accessibility

**Answer: D**

#### Explanation:

To perform maintenance on one of the hosts in a three-node vSAN cluster with the best availability for the VMs with the least effort and data transfer, the maintenance mode option that should be used is Ensure accessibility. This option migrates only enough components to ensure that all accessible VMs remain accessible, but does not guarantee full data redundancy or policy compliance. This option is also the only evacuation mode available for a three-node cluster or a cluster with three fault domains, as there are not enough hosts to perform full data migration or re-protection after a failure. The other options are not correct. Migrating all VMs and their storage from the host to a different storage system or a different vSphere cluster would require more effort and data transfer than using Ensure accessibility, as well as additional resources and configuration steps. Full data migration is not possible in a three-node cluster, as it would require at least four hosts to evacuate all data from one host and maintain full redundancy and policy compliance. References: Place a Member of vSAN Cluster in Maintenance Mode; Working with Maintenance Mode

**NEW QUESTION 5**

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

**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

**NEW QUESTION 6**

A vSAN administrator receives a request from the application team to create a virtual machine on a vSAN datastore. The requirements state that the virtual machine needs to be available quickly after a failure occurs. The solution must minimize administrative effort. Which vSphere feature should the vSAN administrator implement?

- A. Distributed Services Engine
- B. vSphere High Availability
- C. Fault Tolerance
- D. vSphere Distributed Resource Scheduler

**Answer: B**

**Explanation:**

vSphere High Availability is the correct answer because it meets the requirements of making the virtual machine available quickly after a failure occurs and minimizing administrative effort. vSphere HA monitors the health and availability of the hosts and virtual machines in a cluster and automatically restarts any failed virtual machines on other hosts within minutes. vSphere HA also supports proactive HA, which can migrate virtual machines from hosts that are about to fail or have degraded performance. vSphere HA is easy to configure and manage, as it only requires enabling HA on the cluster level and setting some basic policies and options. Distributed Services Engine, Fault Tolerance, and vSphere Distributed Resource Scheduler are not valid or optimal solutions for this scenario. Distributed Services Engine is a new feature in vSphere 7 that provides network services such as firewalling, load balancing, routing, and NAT for virtual machines and containers. It does not directly affect the availability or recovery of virtual machines after a failure. Fault Tolerance provides continuous availability for virtual machines by creating a secondary copy of the virtual machine that runs in lockstep with the primary copy on another host. If the primary copy fails, the secondary copy takes over without any interruption or data loss. However, Fault Tolerance has some limitations and overheads, such as requiring dedicated network bandwidth, supporting only one vCPU per virtual machine, and consuming twice as much CPU and memory resources as a single virtual machine. Fault Tolerance also requires more administrative effort than vSphere HA, as it needs to be enabled and configured for each individual virtual machine. vSphere Distributed Resource Scheduler is a feature that balances the workload and resources across a cluster by automatically migrating virtual machines based on their demand and priority. It does not directly affect the availability or recovery of virtual machines after a failure, although it can work together with vSphere HA to find optimal hosts for restarting failed virtual machines. References:

? [VMware vSAN Specialist v2 Exam Preparation Guide], page 11

? vSphere Availability

? Distributed Services Engine

? vSphere Distributed Resource Scheduler

**NEW 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

**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.

#### NEW QUESTION 8

A customer wants to validate if Skyline online health is working for vSAN and finds out that Skyline is not fully configured yet. What two requirements must be met to make sure that Skyline online health will work? (Choose two.)

- A. Add the Skyline license into Virtual Center
- B. Enable Skyline Health on the vSAN Cluster
- C. Enable CEIP and join the program
- D. Have a working Internet connection
- E. Have vCenter on version 7 or higher

**Answer:** CD

#### Explanation:

To make sure that Skyline online health will work for vSAN, two requirements must be met: enable CEIP and join the program, and have a working Internet connection. CEIP stands for Customer Experience Improvement Program, which is a voluntary program that collects anonymous product usage data from customers who participate in it. By enabling CEIP and joining the program, customers can benefit from Skyline online health, which provides proactive notifications and recommendations for software and hardware issues based on VMware Analytics Cloud. A working Internet connection is also required for Skyline online health to communicate with VMware Analytics Cloud and receive online notifications. The other options are not requirements for Skyline online health. References: About the vSAN Skyline Health; Skyline Health

#### NEW QUESTION 9

An administrator has successfully deployed a vSAN Stretched Cluster and needs to ensure that any virtual machines that are created are placed in the appropriate site. Which two steps are needed to complete this task? (Choose two.)

- A. Create VM/Host groups for the two sites
- B. Create a single VM/Host group across both sites
- C. Put the VMs in a vSphere DRS group
- D. Put the VMs in the correct VM group
- E. Create a storage policy that includes site affinity rules and apply to VMs

**Answer:** AE

#### Explanation:

To ensure that any virtual machines that are created are placed in the appropriate site, the administrator needs to create VM/Host groups for the two sites and create a storage policy that includes site affinity rules and apply to VMs. VM/Host groups allow the administrator to group virtual machines and hosts based on their location or preference. Site affinity rules specify which site a virtual machine should be placed on or prefer to run on. A single VM/Host group across both sites would not allow the administrator to control the placement of virtual machines. Putting the VMs in a vSphere DRS group or in the correct VM group would not affect their site affinity. References: 1, page 12; 2, section 3.2

#### NEW QUESTION 10

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.

**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]

#### NEW QUESTION 10

An administrator is tasked to create a Kerberos secured NFS v4.1 file share. Which information is minimally required during the configuration of the File Service?



- A. Organizational Unit, User Account, Password
- B. Active Directory Domain, User Account, Password
- C. Kerberos Server, User Account, Password
- D. Active Directory Domai
- E. Organizational Unit, User Accoun
- F. Password

**Answer: B**

**Explanation:**

To create a Kerberos secured NFS v4.1 file share, the administrator needs to provide the following information during the configuration of the File Service:  
 ? Active Directory Domain: The domain name of the Active Directory server that provides Kerberos authentication service for the NFS server and clients. For example, example.com.  
 ? User Account: The user name of the Active Directory account that has permissions to join the NFS server to the domain and create service principal names (SPNs) for the NFS server. For example, administrator@example.com.  
 ? Password: The password of the Active Directory account that is used for authentication. For example, P@ssw0rd.  
 These information are required to enable Kerberos security for NFS 4.1 and allow the NFS server to obtain a Kerberos ticket from the Active Directory server. The administrator also needs to specify the NFS share name, path, and access permissions1 References: 1: VMware vSphere Storage Guide, page 118

**NEW QUESTION 11**

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. Tanzu Kubernetes Grid Integrated Edition

**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

**NEW QUESTION 14**

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.

**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

**NEW QUESTION 19**

Refer to the exhibit.

An administrator uses SSH to log into a vSAN ESA host and runs theesxcli vsan debug object overviewcommand.

Object GUID	Group GUID	Version	Size	Used	SPBM Profile	Healthy Components
19413f63-84bd-4aba-2ba6-0050560659c0	1a413f63-a8d1-fafb-0809-0050560659c0	17	0.12 GB	0.01 GB	vSAN Default Storage Policy	7 of 8
1c413f63-4c1a-73bc-9046-0050560659c0	1a413f63-a8d1-fafb-0809-0050560659c0	17	1.00 GB	0.70 GB	vSAN Default Storage Policy	5 of 8
1a413f63-a8d1-fafb-0809-0050560659c0	1a413f63-a8d1-fafb-0809-0050560659c0	17	255.00 GB	0.05 GB	vSAN Default Storage Policy	4 of 8
cf403f63-eec5-da41-8599-005056065997	cf403f63-eec5-da41-8599-005056065997	17	255.00 GB	0.04 GB	vSAN Default Storage Policy	7 of 8
d0403f63-f7af-45cd-1e8a-005056065997	cf403f63-eec5-da41-8599-005056065997	17	1.00 GB	0.70 GB	vSAN Default Storage Policy	7 of 8
ef403f63-fe7b-46f0-9d4f-005056065997	cf403f63-eec5-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	1.00 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
0d413f63-8c5e-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.

**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

**NEW QUESTION 23**

After reviewing various performance charts at a cluster level, an administrator found an individual VM impacting overall performance of the vSAN cluster. What feature should be used to introspect multiple performance metrics of a single virtual machine?

- A. esxci
- B. Skyline Health
- C. I/O Trip Analyzer
- D. IIOInsight

**Answer:** C

**Explanation:**

To introspect multiple performance metrics of a single virtual machine, such as latency, throughput, IOPS, and congestion, the feature that should be used is I/O Trip Analyzer. This feature allows the administrator to diagnose the virtual machine I/O latency issues by providing a breakdown of the latencies at each layer of the vSAN stack, such as VM, host, network, and disk group. The other options are not correct, as they do not provide multiple performance metrics of a single virtual machine. esxcli is a command-line tool that can be used to manage various aspects of ESXi hosts, but it does not provide detailed performance analysis of virtual machines. Skyline Health is a feature that provides proactive notifications and recommendations for software and hardware issues based on VMware Analytics Cloud, but it does not provide granular performance metrics of virtual machines. IIOInsight is not a valid feature name in vSAN. References: Use I/O Trip Analyzer; Monitoring vSAN Performance

**NEW QUESTION 26**

A vSAN administrator is tasked to perform an upgrade of a vSAN cluster, including firmware and drivers for its hardware. The vSAN administrator already created an image using vSphere Lifecycle Manager (vLCM).

Prior to selecting Start Remediation, which step should be taken to upgrade the complete vSAN cluster as a single task?

- A. Select Remediate All through vLCM to upgrade all hosts in the cluster
- B. Place all hosts in the vSAN cluster into Maintenance Mode
- C. Stage the upgrade of the vSAN cluster through vLCM
- D. Manually remediate one host at a time in the vSAN cluster

**Answer:** A

**Explanation:**

To upgrade the complete vSAN cluster as a single task, including firmware and drivers for its hardware, the vSAN administrator should select Remediate All through vLCM to upgrade all hosts in the cluster. This option allows the administrator to apply the image created by vLCM to all hosts in the cluster in a single operation, without having to manually remediate each host individually. The other options are not correct, as they do not perform the upgrade of the vSAN cluster as a single task. Placing all hosts in the vSAN cluster into Maintenance Mode is not necessary, as vLCM will automatically place each host into Maintenance Mode before applying the image. Staging the upgrade of the vSAN cluster through vLCM is only a preparatory step that downloads the image components to each host, but does not apply them. Manually remediating one host at a time in the vSAN cluster is not efficient, as it requires more user intervention and time. References: vSphere Lifecycle Manager (vLCM) on HPE; Lifecycle Management with vLCM in vSAN 7 Update 1

**NEW QUESTION 28**

After a planned power outage, an administrator decided to restart the vSAN cluster manually.

What is the correct sequence of steps for the administrator to follow after powering on the ESXi hosts?

- A. \* 1. Enable cluster member updates from vCenter Server only on one ESXi host.\* 2. Run the python reboot helper script only on one ESXi host to recover the cluster.\* 3. Exit all hosts from maintenance mode.
- B. \* 1. Exit all hosts from maintenance mode.\* 2. Run the python reboot helper script only on one ESXi host to recover the cluster.\* 3. Enable cluster member updates from vCenter Server on all ESXi hosts.
- C. \* 1. Exit all hosts from maintenance mode.\* 2. Enable cluster member updates from vCenter Server only on one ESXi host.\* 3. Run the python reboot helper script only on one ESXi host to recover the cluster.
- D. \* 1. Enable cluster member updates from vCenter Server on all ESXi hosts.\* 2. Run the python reboot helper script on all ESXi hosts to recover the cluster.\* 3. Exit all hosts from maintenance mode.

**Answer:** A

**Explanation:**

This is the sequence of steps recommended by VMware for manually restarting the vSAN cluster after a planned power outage. The steps are as follows:

? Enable cluster member updates from vCenter Server only on one ESXi host. This will allow the host to receive the latest cluster membership information from vCenter Server and avoid any conflicts or inconsistencies with other hosts. The command to enable cluster member updates is `esxcfg-advcfg -s 1 /VSAN/IgnoreClusterMemberListUpdates`.

? Run the python reboot helper script only on one ESXi host to recover the cluster.

This will prepare the cluster for a manual restart by partitioning the cluster and ensuring that all hosts have consistent metadata. The command to run the python reboot helper script is `python /usr/lib/vmware/vsan/bin/reboot_helper.py prepare`.

? Exit all hosts from maintenance mode. This will allow the hosts to resume normal operations and join the vSAN cluster. The command to exit maintenance mode is `esxcli system maintenanceMode set -e false`.

The other options are incorrect for the following reasons:

? B, exit all hosts from maintenance mode, run the python reboot helper script only on one ESXi host to recover the cluster, and enable cluster member updates from vCenter Server on all ESXi hosts, is incorrect because exiting all hosts from maintenance mode before running the python reboot helper script can cause data inconsistency or corruption, as the hosts may not have the latest metadata or cluster membership information. Enabling cluster member updates from vCenter

Server on all ESXi hosts is also unnecessary and can cause conflicts or inconsistencies with other hosts.

? C, exit all hosts from maintenance mode, enable cluster member updates from vCenter Server only on one ESXi host, and run the python reboot helper script only on one ESXi host to recover the cluster, is incorrect because exiting all hosts from maintenance mode before running the python reboot helper script can cause data inconsistency or corruption, as the hosts may not have the latest metadata or cluster membership information.

? D, enable cluster member updates from vCenter Server on all ESXi hosts, run the python reboot helper script on all ESXi hosts to recover the cluster, and exit all hosts from maintenance mode, is incorrect because enabling cluster member updates from vCenter Server on all ESXi hosts is unnecessary and can cause conflicts or inconsistencies with other hosts. Running the python reboot helper script on all ESXi hosts concurrently can also cause a race condition that can result in unexpected outcomes.

References:

? Manually Shut Down and Restart the vSAN Cluster

? Restart the vSAN Cluster

### NEW QUESTION 32

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

**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

### NEW QUESTION 33

A vSAN administrator is responsible for managing a customer's production vSAN cluster that is going to be used to provide SMB file shares to a number of host clients. The vSAN administrator must take action so the performance of all services in the production vSAN cluster can be monitored.

Which two services must be enabled for this monitoring to occur? (Choose two.)

- A. vSAN Performance Diagnostic Service
- B. iSCSI Target Service
- C. vSAN File Services
- D. vSAN Health Service
- E. vSAN Performance Service

**Answer:** CE

#### Explanation:

To monitor the performance of vSAN File Services, the vSAN administrator must enable both the vSAN File Services and the vSAN Performance Service. The vSAN File Services provides SMB file shares to host clients, while the vSAN Performance Service collects and analyzes performance statistics and displays them in the vSphere Client. The other services are not related to vSAN File Services performance monitoring.

References: VMware vSAN Specialist v2 EXAM 5V0-22.23, page 9, Objective 7.4; [vSAN File Services]; [vSAN Performance Service]

### NEW QUESTION 36

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

**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

### NEW QUESTION 38

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

**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

**NEW QUESTION 40**

An existing vSAN OSA cluster has this specification: Four ESXi hosts with all flash configuration

Each with two disk groups

Each disk group with one cache device and four capacity devices There are five more device slots available per host

The CTO would like to provision new applications, and these will need more capacity and performance.

Which two methods should be used by the vSAN administrator to meet this goal with the least amount of impact? (Choose two.)

- A. Replacing all capacity devices with a similar larger device
- B. Replacing all cache devices with a larger device
- C. Adding one more disk group per host with the same configuration
- D. Adding faster cache devices
- E. Adding an ESXi host with identical device configuration

**Answer:** CE

**Explanation:**

Adding one more disk group per host with the same configuration and adding an ESXi host with identical device configuration are the two methods that the vSAN administrator should use to meet the goal of increasing capacity and performance with the least amount of impact. Adding one more disk group per host will increase the raw storage capacity by 20% and also improve the performance by distributing the I/O load across more cache devices and disk groups. Adding an ESXi host with identical device configuration will increase the raw storage capacity by 25% and also improve the performance by adding more compute and network resources to the cluster. Both methods can be done without disrupting any ongoing operations or requiring any data evacuation or resynchronization. The other options are incorrect for the following reasons:

? Replacing all capacity devices with a similar larger device is incorrect because it will not increase the performance and will have a significant impact on the cluster. Replacing the capacity devices requires deleting the disk groups, which will erase all data on them and trigger a resynchronization of the affected objects. This can be disruptive and time-consuming, and also introduce additional network and disk traffic.

? Replacing all cache devices with a larger device is incorrect because it will not increase the capacity and will have a significant impact on the cluster. Replacing the cache devices also requires deleting the disk groups, which will have the same drawbacks as replacing the capacity devices. Moreover, increasing the cache size may not improve the performance significantly, as vSAN OSA uses a fixed cache ratio of 70% for write buffer and 30% for read cache, regardless of the cache device size.

? Adding faster cache devices is incorrect because it will not increase the capacity and will have a significant impact on the cluster. Adding faster cache devices also requires deleting the disk groups, which will have the same drawbacks as replacing the cache devices. Furthermore, adding faster cache devices may not improve the performance significantly, as vSAN OSA uses a fixed cache ratio of 70% for write buffer and 30% for read cache, regardless of the cache device speed. References:

? VMware vSAN Specialist v2 Exam Preparation Guide, page 10

? Expanding a vSAN Cluster

**NEW QUESTION 45**

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

**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

**NEW QUESTION 49**

.....



## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### 5V0-22.23 Practice Exam Features:

- \* 5V0-22.23 Questions and Answers Updated Frequently
- \* 5V0-22.23 Practice Questions Verified by Expert Senior Certified Staff
- \* 5V0-22.23 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* 5V0-22.23 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The 5V0-22.23 Practice Test Here](#)**