

# 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 vSAN 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. 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

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

**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

#### NEW QUESTION 5

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

**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

**NEW QUESTION 6**

What are two prerequisites for using the TRIM and UNMAP capability of vSAN? (Choose two.)

- A. Deduplication and compression are enabled.
- B. The vSAN cluster is an all-flash architecture.
- C. The VM guest operating system supports ATA TRIM or SCSI UNMAP capability
- D. TRIM and UNMAP is enabled.
- E. Change the Object Space Reservation to 100.

**Answer: BD**

**Explanation:**

The two prerequisites for using the TRIM and UNMAP capability of vSAN are:

? B. The vSAN cluster is an all-flash architecture. TRIM and UNMAP are only supported on all-flash vSAN clusters, as they can reclaim space from flash devices that use thin provisioning. TRIM and UNMAP are not supported on hybrid vSAN clusters, as they cannot reclaim space from magnetic disks that use thick provisioning1.

? D. TRIM and UNMAP is enabled. TRIM and UNMAP are disabled by default in vSAN, as they might have a performance impact on some workloads. To enable TRIM and UNMAP on a vSAN cluster, the administrator must use the following RVC command: `vsan.unmap_support -enable2`. After enabling TRIM and UNMAP, the administrator must power off and then power on all VMs that use the vSAN datastore.

**NEW QUESTION 7**

An administrator has deployed a new vSAN OSA cluster that contains eight hosts and needs to configure a storage policy for the currently deployed database virtual machines. The requirements state that if two hosts in the vSAN OSA cluster fail, all virtual machines are unaffected.

Which RAID configuration must the administrator use in this storage policy to achieve the best performance for the database virtual machines?

- A. RAID-1
- B. RAID-5
- C. RAID-6
- D. RAID-0

**Answer: A**

**Explanation:**

To achieve the best performance for the database virtual machines and tolerate two host failures in a vSAN OSA cluster, the administrator must use RAID-1 as the RAID configuration in the storage policy. RAID-1 is a mirroring technique that creates multiple replicas of each object across different hosts. RAID-1 provides the best performance among the available RAID configurations, as it does not involve any parity calculations or stripe splitting. To tolerate two host failures, the administrator must set the Failures to Tolerate (FTT) policy to 2, which means that each object will have three replicas. The other options are not correct. RAID-5 and RAID-6 are erasure coding techniques that split each object into data segments and parity segments across different hosts. RAID-5 can tolerate one host failure, while RAID-6 can tolerate two host failures. However, both RAID-5 and RAID-6 have lower performance than RAID-1, as they involve more complex calculations and network traffic. RAID-0 is a striping technique that splits each object into multiple stripes across different hosts. RAID-0 does not provide any data redundancy or fault tolerance, and therefore cannot tolerate any host failure.

References: RAID Configurations, FTT, and Host Requirements; RAID 5 or RAID 6 Design Considerations

**NEW QUESTION 8**

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 9

How often does the Skyline Health interval validate online if there are new Health Checks available for vSAN?

- A. Every 1 hour
- B. Every 4 hours
- C. Every 24 hours
- D. Every 12 hours

**Answer: C**

#### Explanation:

The Skyline Health interval validates online if there are new Health Checks available for vSAN every 24 hours. This means that vSAN checks for new health checks from VMware Analytics Cloud once a day and updates the vSAN Health Service accordingly. The other options are not correct, as they do not match the actual frequency of the online validation. References: About the vSAN Skyline Health

#### NEW QUESTION 10

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.

**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

#### NEW QUESTION 10

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

**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

#### NEW QUESTION 11

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

**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  
 12 References: 1: VMware vSAN Specialist v2 Exam Preparation Guide, page 14 2: VMware vSAN Design and Sizing Guide, page 32

#### NEW QUESTION 12

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 17

An administrator has been tasked with upgrading existing vSAN OSA cluster hosts with a SSD cache device per host to a NVMe device (hot plug). Which fact should guide the administrator's action?

- A. The disk group must be deleted on each physical host in the vSAN OSA cluster to use the NVMe device.
- B. The disk group does not need to be removed before adding new cache.
- C. The host must be removed from vSAN OSA cluster before changing cache devices.
- D. The cache disk drives must have a larger capacity.

**Answer: A**

#### Explanation:

The correct answer is A, the disk group must be deleted on each physical host in the vSAN OSA cluster to use the NVMe device. This is because vSAN OSA uses a disk group configuration where one disk is designated as a cache disk and the rest are capacity disks. To replace the cache disk with a different type or size, the disk group must be deleted first, which will erase all data on the disks and trigger a resynchronization of the affected objects. The administrator should put the host in maintenance mode and choose the option to evacuate all data before deleting the disk group. After replacing the cache disk with the NVMe device, the administrator should recreate the disk group and exit maintenance mode. The other options are incorrect for the following reasons:

? B, the disk group does not need to be removed before adding new cache, is incorrect because adding a new cache disk to an existing disk group is not supported in vSAN OSA. The cache disk can only be replaced by deleting and recreating the disk group.

? C, the host must be removed from vSAN OSA cluster before changing cache devices, is incorrect because removing the host from the cluster is not necessary and will cause more disruption and data loss than putting the host in maintenance mode. Removing the host will also delete its disk groups and require re-adding them after rejoining the cluster.

? D, the cache disk drives must have a larger capacity, is incorrect because there is no requirement for the cache disk to have a larger capacity than the existing one. The cache disk size should be determined by the workload characteristics and performance requirements, not by the expansion process. References:

? VMware vSAN Specialist v2 Exam Preparation Guide, page 10

#### NEW QUESTION 19

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 22

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 23

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 26

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

**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]

#### NEW QUESTION 29

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 30

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

**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

### NEW QUESTION 31

A vSAN administrator has recently upgraded a vSAN cluster to 8.0 OSA and has enabled Capacity Reserve features to reduce the amount of capacity reserved for transient and rebuild operations.

Which scenario would prevent this feature from operating properly?

- A. Underutilized space is above 25-30% of the total capacity threshold.
- B. The used space on vSAN datastore exceeds the suggested host rebuild threshold.
- C. The used space on vSAN datastore exceeds the suggested slack rebuild threshold.
- D. The physical disk has reached an 80% full reactive rebalance threshold.

**Answer:** C

#### Explanation:

The Capacity Reserve feature in vSAN 8.0 OSA reduces the amount of capacity reserved for transient and rebuild operations by using a slack space threshold. This threshold is calculated based on the size of the largest component in the cluster and the number of failures to tolerate. If the used space on vSAN datastore exceeds the suggested slack space threshold, the feature will not operate properly and vSAN will revert to using the host rebuild reserve threshold. The other scenarios will not affect the Capacity Reserve feature. References: [VMware vSAN Specialist v2 EXAM 5V0-22.23], page 28

### NEW QUESTION 35

An administrator wishes to prevent vCenter notifications of vSAN Health status during a scheduled maintenance window.

Which action can be taken to achieve this goal?

- A. Disable the performance service
- B. Disable the alarm from the HTML client
- C. Run performance diagnostics prior to scheduled maintenance
- D. Disable SNMP service

**Answer:** B

#### Explanation:

To prevent vCenter notifications of vSAN Health status during a scheduled maintenance window, the administrator can disable the alarm from the HTML client. This will suppress the alerts for a specified duration or until the alarm is re-enabled. Disabling the performance service, running performance diagnostics, or disabling SNMP service will not affect the vSAN Health status notifications. References: [VMware vSAN Specialist v2 EXAM 5V0-22.23], page 25

### NEW QUESTION 39

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

**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

### NEW QUESTION 42

A vSAN administrator notices that the VMware Skyline Health: Network Latency Check reports indicate that three hosts are noncompliant.

Which action should the vSAN administrator take?

- A. Immediately reboot the non-compliant hosts
- B. Check VMKNICs, uplinks, physical switches, and associated settings
- C. Rerun the VMware Skyline Health: vSAN Cluster Partition report
- D. Place the non-compliant hosts into an isolated network partition

**Answer: B**

#### Explanation:

The correct answer is B, check VMKNICs, uplinks, physical switches, and associated settings. This is because the VMware Skyline Health: Network Latency Check reports the network latency between vSAN hosts and displays the network latency in real time. Failure indicates that the network latency is above the normal threshold, which can affect the performance and availability of vSAN. The network latency can be caused by various factors, such as misconfiguration, congestion, or errors in the network components. The vSAN administrator should check the VMKNICs, uplinks, physical switches, and associated settings for any issues and resolve them accordingly. The vSAN administrator can also use tools such as vmkping or esxtop to test the network connectivity and performance between hosts. The other options are incorrect for the following reasons:

? A, immediately reboot the non-compliant hosts, is incorrect because rebooting the non-compliant hosts is not a recommended action and can cause more disruption and data loss than resolving the network issue. Rebooting the hosts will also trigger a resynchronization of data across the cluster, which can affect the performance and availability of vSAN.

? C, rerun the VMware Skyline Health: vSAN Cluster Partition report, is incorrect because rerunning the VMware Skyline Health: vSAN Cluster Partition report will not help to resolve the network latency issue. The vSAN Cluster Partition report checks if there are any network partitions in the cluster that prevent communication between hosts. The network partition can be caused by network latency, but it is not the same as network latency. The vSAN administrator should first fix the network latency issue before checking for any network partitions.

? D, place the non-compliant hosts into an isolated network partition, is incorrect because placing the non-compliant hosts into an isolated network partition will not help to resolve the network latency issue. It will also cause more problems for vSAN, such as data inconsistency, reduced redundancy, and degraded performance.

The vSAN administrator should avoid creating any network partitions in the cluster and ensure that all hosts can communicate with each other. References:

? VMware vSAN Specialist v2 Exam Preparation Guide, page 9

? Network Health - Network Latency Check (2149511)

### NEW QUESTION 45

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 48

Which two actions are recommended when adding a host to a vSAN cluster? (Choose two.)

- A. Create uniformly-configured hosts

- B. Disable vSAN performance service
- C. Disable vSphere Cluster Services
- D. Disable vSphere High Availability (HA)
- E. Reference the VMware Compatibility Guide

**Answer:** AE

**Explanation:**

When adding a host to a vSAN cluster, it is recommended to create uniformly-configured hosts and reference the VMware Compatibility Guide. These actions will ensure that the host meets the hardware and software requirements for vSAN, and that it can work seamlessly with the existing hosts in the cluster. Uniformly-configured hosts have the same number and type of disk groups, cache devices, capacity devices, network adapters, and drivers. The VMware Compatibility Guide provides a list of certified components and firmware versions that are compatible with vSAN. The other options are not recommended, as they can cause disruption or degradation of the vSAN cluster services. Disabling vSAN performance service, vSphere Cluster Services, or vSphere High Availability (HA) can affect the monitoring, availability, and load balancing of the cluster.

**NEW QUESTION 50**

What are two characteristics of a durability component in vSAN? (Choose two.)

- A. Better Performance
- B. Faster resynchronization
- C. Faster snapshot creation
- D. Better Storage utilization
- E. Better Availability

**Answer:** BE

**Explanation:**

A durability component is a temporary component that is created when a host or disk group is placed in maintenance mode with the Ensure data accessibility option, or when a host or disk group fails unexpectedly. A durability component improves the availability of data by maintaining the required number of failures to tolerate (FTT) until the original component is restored or rebuilt. A durability component also speeds up the resynchronization process by reducing the amount of data that needs to be copied. The other characteristics are not applicable to a durability component. References: VMware vSAN Specialist v2 EXAM 5V0-22.23, page 10, Objective 6.8; [Durability Components]

**NEW QUESTION 52**

An administrator is upgrading multiple vSAN Witness nodes with vLCM (single image management) that are used for vSAN Stretched and two-node Clusters. What two witness node types can the administrator upgrade? (Choose two.)

- A. Appliance witness node
- B. Shared witness node
- C. Nested witness node
- D. Dedicated witness node
- E. Physical witness node

**Answer:** AC

**Explanation:**

To upgrade multiple vSAN Witness nodes with vLCM (single image management) that are used for vSAN Stretched and two-node Clusters, the administrator can upgrade two witness node types: appliance witness node and nested witness node. An appliance witness node is a virtual ESXi host that runs on a physical ESXi host and contains the witness components of VM objects stored in the vSAN cluster. A nested witness node is a virtual ESXi host that runs on another virtual ESXi host and contains the witness components of VM objects stored in the vSAN cluster. Both types of witness nodes can be managed by vLCM as independent nodes since vSphere 7.0 Update 3, as long as they are version 7.0 Update 2 or later. The other options are not correct. A shared witness node is a witness node that serves multiple vSAN clusters, which is not supported by vLCM. A dedicated witness node is a witness node that serves only one vSAN cluster, which is not a specific type of witness node. A physical witness node is a physical ESXi host that contains the witness components of VM objects stored in the vSAN cluster, which cannot be upgraded by vLCM. References: vSphere Lifecycle Manager and the vSAN Witness Hosts; Shared Witness for 2-Node vSAN Deployments

**NEW QUESTION 53**

A vSAN administrator of a non-internet connected vSAN environment wants to upgrade the environment from the vSAN 7.0 U3 to the vSAN 8.0 using vLCM. Which option, if any, should be used as a depot in this case?

- A. Configure the vSphere Lifecycle Manager to download the updates from an Online Depot.
- B. Configure the vSphere Lifecycle Manager to download the updates from the VMware Depot using HTTPS.
- C. Configure the vSphere Lifecycle Manager to download updates from a local UMDS- shared repository.
- D. It is not possible to use the vSphere Lifecycle Manager on a non-internet connected environment.

**Answer:** C

**Explanation:**

To upgrade the vSAN environment from vSAN 7.0 U3 to vSAN 8.0 using vLCM in a non-internet connected environment, the administrator should configure the vSphere Lifecycle Manager to download updates from a local UMDS-shared repository. UMDS stands for Update Manager Download Service, which is a component of vSphere Lifecycle Manager that can be used to download patches and updates for ESXi hosts, virtual appliances, and VMware Tools from the VMware online depot and store them in a shared repository. The administrator can then configure the vSphere Lifecycle Manager to use the UMDS-shared repository as a custom depot for patching and upgrading the vSAN cluster. This option allows the administrator to perform offline upgrades without requiring internet access for the vSAN cluster. References: 1: VMware vSphere Lifecycle Manager Administration, page 22 2: VMware vSphere Update Manager Download Service, page 5

**NEW QUESTION 58**

A vSAN administrator has an existing cluster where each ESXi host has the following: Disk group #1 with one cache device and three capacity devices. Disk group #2 with one cache device and two capacity devices.

What must the vSAN administrator do to expand disk group #2 to have three capacity devices?

- A. Create a new disk group with a single capacity device and then migrate the existing capacity devices
- B. Add the new capacity device to the disk group and vSAN will automatically rebalance
- C. Put the entire ESXi host in maintenance mode, evacuate all data, then add the new capacity device
- D. Put the disk group in maintenance mode, evacuate all data, then add the new capacity device

**Answer: B**

**Explanation:**

To expand disk group #2 to have three capacity devices, the vSAN administrator should add the new capacity device to the disk group and vSAN will automatically rebalance. This action allows the administrator to increase the storage capacity of the disk group without disrupting any ongoing operations or evacuating any data. vSAN will automatically distribute data across all devices in the disk group to balance performance and utilization. The other options are not correct. Creating a new disk group with a single capacity device and then migrating the existing capacity devices is not necessary, as it would require more steps and resources than adding a device to an existing disk group. Putting the entire ESXi host or the disk group in maintenance mode and evacuating all data is not required, as it would cause downtime and data movement that could be avoided by adding a device to an existing disk group. References: Add Devices to the Disk Group; Expanding a vSAN Cluster

**NEW QUESTION 59**

Due to a planned power outage, an administrator decides to shut down the vSAN cluster using the Shutdown Cluster Wizard. The administrator starts by checking the vSAN health service to confirm the cluster is healthy and then powers off all virtual machines (VMs) including vCLS VMs.

Which step needs to be taken before starting the Shutdown Cluster Wizard?

- A. Place all ESXi hosts into maintenance mode
- B. Disable cluster member updates from vCenter Server
- C. Turn off High Availability
- D. Shutdown vCenter

**Answer: C**

**Explanation:**

To shut down the vSAN cluster using the Shutdown Cluster Wizard, the administrator needs to turn off High Availability (HA) before starting the wizard. This is because HA monitors the cluster for host failures and attempts to restart the affected VMs on other hosts. If HA is not turned off, the cluster might register host shutdowns as failures and trigger unnecessary VM restarts, which can interfere with the graceful shutdown process. Therefore, the administrator should disable HA from the Configure tab of the cluster before using the Shutdown Cluster Wizard. References: 1: Shut Down the vSAN Cluster Using the Shutdown Cluster Wizard 3 2: Manually Shut Down and Restart the vSAN Cluster 4

**NEW QUESTION 61**

Which two considerations should an architect assess when designing a HCI Mesh solution with VMware vSAN and VMware vSphere High Availability (HA)? (Choose two.)

- A. A server vSAN cluster can serve its local datastore up to five client vSAN clusters.
- B. A client cluster can mount up to ten remote datastores from one or more vSAN server clusters.
- C. A minimum of three nodes are required within the client cluster for vSphere HA to work
- D. If vSphere HA is to work with HCI Mesh, Datastore with Permanent Device Loss (PDL) must be configured to Power off and restart VMs.
- E. If vSphere HA is to work with HCI Mesh, Datastore with AllPaths Down (APD) must be configured to Power off and restart VMs.

**Answer: CE**

**Explanation:**

To design a HCI Mesh solution with VMware vSAN and VMware vSphere High Availability (HA), two considerations that the architect should assess are: A minimum of three nodes are required within the client cluster for vSphere HA to work. This is because vSphere HA needs at least three nodes in a cluster to form a quorum and elect a master host that monitors the availability of other hosts and VMs. If there are less than three nodes in a cluster, vSphere HA cannot function properly and might fail to detect or respond to host or VM failures. If vSphere HA is to work with HCI Mesh, Datastore with All Paths Down (APD) must be configured to Power off and restart VMs. This is because APD is a condition that occurs when a storage device becomes inaccessible due to loss of physical connectivity, resulting in I/O errors or timeouts for VMs that use that device. When using HCI Mesh, APD can happen if the network connection between the client cluster and the server cluster is lost or disrupted, causing the remote datastore to become unavailable. To ensure that vSphere HA can restart the affected VMs on another host that has access to their storage, Datastore with APD must be set to Power off and restart VMs in the vSphere HA settings. The other options are not correct. A server vSAN cluster can serve its local datastore up to 15 client vSAN clusters, not five. This is the maximum number of client clusters that can mount a remote datastore from a server cluster using HCI Mesh. A client cluster can mount up to five remote datastores from one or more vSAN server clusters, not ten. This is the maximum number of remote datastores that can be mounted by a client cluster using HCI Mesh. References: VMware vSAN HCI Mesh; vSphere Availability; Handling All Paths Down (APD) Conditions

**NEW QUESTION 65**

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

**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

#### NEW QUESTION 67

A vSAN administrator wants to transition from VMware Update Manager to vSphere Lifecycle Manager. Which element is a mandatory requirement to create an image?

- A. ESXi Version
- B. Component
- C. Firmware and Drivers Add-On
- D. Vendor Add-On

**Answer:** A

#### Explanation:

To create an image using vSphere Lifecycle Manager, the mandatory requirement is to specify the ESXi version. An image is a collection of software components that define the desired state of hosts in a cluster. An image must include at least one ESXi version component, which determines the base hypervisor software for the hosts. Optionally, an image can also include other components, such as vendor add-ons, firmware and drivers add-ons, or custom components. The other options are not correct. A component is a generic term for any software element that can be included in an image, but it is not a specific type of component. A firmware and drivers add-on is an optional component that provides firmware and drivers updates for hardware devices on the hosts. A vendor add-on is an optional component that provides vendor-specific software for the hosts. References: About Images; Create an Image

#### NEW QUESTION 70

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

**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

#### NEW QUESTION 74

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

**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

#### NEW QUESTION 79

.....

## **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](#)**