

Exam Questions 5V0-22.23

VMware vSAN Specialist (v2)

<https://www.2passeasy.com/dumps/5V0-22.23/>



NEW QUESTION 1

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 2

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 3

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 4

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 5

A vSAN administrator is noticing that the objects resynchronizing in the cluster are taking longer than expected and wants to view the resynchronizing metrics. Which performance category should the vSAN administrator open?

- A. Disks
- B. Host Network
- C. Resvnc Latency
- D. Backend

Answer: D

Explanation:

To view the resynchronizing metrics, the vSAN administrator should open the Backend performance category. This category shows the performance of vSAN data components, such as read/write latency, IOPS, throughput, congestion, and resync traffic. The other categories are not relevant for this task. Disks shows the performance of physical disks in the cluster, Host Network shows the network performance of vSAN hosts, and Resvnc Latency shows the latency of resynchronization operations. References: 1, page 23; 3, section 6.4

NEW QUESTION 6

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

Which set of steps should the vSAN administrator take?

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

Answer: C

Explanation:

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

NEW QUESTION 7

An administrator must choose between deploying a virtual witness or a physical witness for a vSAN Stretched Cluster. The administrator eventually decides to use a virtual witness.

What is a benefit of selecting this approach?

- A. Increased vSAN datastore capacity
- B. Shared metadata between separate clusters
- C. Reduced vSphere licensing
- D. Additional compute capacity for running VMs

Answer: C

Explanation:

The correct answer is C, reduced vSphere licensing. This is because using a virtual witness appliance instead of a physical witness host can save on vSphere licensing costs, as the virtual witness appliance does not consume a vSphere license. The virtual witness appliance is a preconfigured virtual machine that runs ESXi and is distributed as an OVA file. It can be deployed on any ESXi host that has network connectivity to both data sites of the stretched cluster. The virtual witness appliance does not run any virtual machines other than itself and only hosts witness components of virtual machine objects. The other options are incorrect for the following reasons:

? A, increased vSAN datastore capacity, is incorrect because using a virtual witness appliance does not affect the vSAN datastore capacity. The witness appliance does not store any customer data, only metadata, such as the size and UUID of vSAN object and components. The witness appliance also does not contribute any storage devices to the vSAN datastore.

? B, shared metadata between separate clusters, is incorrect because using a virtual witness appliance does not enable sharing metadata between separate clusters. The witness appliance is dedicated to one stretched cluster and cannot serve as a witness for multiple clusters. The witness appliance maintains consistency between the two data sites of the stretched cluster by hosting witness components that act as tie-breakers in case of a site failure or network partition.

? D, additional compute capacity for running VMs, is incorrect because using a virtual witness appliance does not provide additional compute capacity for running VMs. The witness appliance does not run any VMs other than itself and does not participate in any compute operations of the stretched cluster. The witness appliance only hosts witness components that consume minimal CPU and memory resources. References:

? VMware vSAN Specialist v2 Exam Preparation Guide, page 11

? Deploying a vSAN Witness Appliance

NEW QUESTION 8

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 9

The vSphere Client reports that the state of some components stored on the vSAN datastore are in the reconfiguring state. Which situation causes components to enter this state?

- A. A host in the cluster enters maintenance mode.
- B. The cluster is recovering from a vSAN failure.
- C. The applied storage policy is modified.
- D. Additional storage capacity is added to the cluster.

Answer: C

Explanation:

The reconfiguring state indicates that some components stored on the vSAN datastore are being moved or resized to meet a new storage policy requirement. This state can occur when the applied storage policy is modified, such as changing the number of failures to tolerate, stripe width, or object space reservation. The other situations will not cause components to enter this state. References: [VMware vSAN Specialist v2 EXAM 5V0-22.23], page 31

NEW QUESTION 10

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

- A. vSphere -> vCenter -> vSAN on-disk format
- B. vSphere -> vSAN on-disk format -> vCenter
- C. vCenter -> vSphere -> vSAN on-disk format
- D. vSAN on-disk format -> vSphere -> vCenter

Answer: C

Explanation:

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

NEW QUESTION 10

All of the virtual machines running on a hybrid vSAN datastore have this storage policy assigned:

Failures to Tolerate (FTT) rule is set to "2 Failures - RAID-1 (Mirroring)"

The vSAN administrator needs to reduce the amount of vSAN datastore capacity the virtual machines will consume.

Which action should the vSAN administrator take to meet this goal?

- A. Modify the FTT rule to "2 Failures - RAID-5 (Erasure Coding)"
- B. Add the "Flash read cache reservation" rule to the storage policy, and set to 0%
- C. Disable Operations reserve and Host rebuild reserve and click "Apply"
- D. Change the FTT rule to "1 Failure - RAID-1 (Mirroring)", and select "Now" for Reapply to VMs

Answer: D

Explanation:

To reduce the amount of vSAN datastore capacity the virtual machines will consume, the vSAN administrator should change the FTT rule to "1 Failure - RAID-1 (Mirroring)", and select "Now" for Reapply to VMs. This action will reduce the number of replicas for each object from three to two, and thus free up some space on the vSAN datastore. The other options are not correct, as they will not reduce the capacity consumption. Modifying the FTT rule to "2 Failures - RAID-5 (Erasure Coding)" will not work for a hybrid vSAN cluster, as erasure coding is only supported for all-flash clusters. Adding the "Flash read cache reservation" rule to the storage policy, and setting to 0% will not affect the capacity layer, as it only controls the amount of flash cache reserved for each object. Disabling Operations reserve and Host rebuild reserve and clicking "Apply" will not change the actual space used by the objects, as these reserves are only logical settings that affect how much free space is reported by vSAN. References: 1, page 9; , section 4.3

NEW QUESTION 15

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

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

Which vSAN feature should be used for this purpose?

- A. vSAN File Services
- B. vSAN HCI Mesh
- C. vSAN Replication

D. vSAN Stretched Clusters

Answer: B

Explanation:

To allocate additional storage from another vSAN enabled vSphere cluster to this vSAN enabled Cluster, the administrator should use the vSAN HCI Mesh feature. This feature allows a vSAN cluster to consume storage resources from another vSAN cluster without requiring the hosts to be part of the same cluster. This way, the administrator can leverage the unused or underutilized storage capacity from another cluster and avoid purchasing new hardware or migrating VMs. The vSAN HCI Mesh feature also supports storage policies, encryption, deduplication and compression, and erasure coding across clusters. References: 1: VMware vSAN Specialist v2 Exam Preparation Guide, page 15 2: VMware vSAN 7 Update 1 - HCI Mesh 3

NEW QUESTION 16

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 20

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 25

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

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

Answer: A

Explanation:

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

References: Introduction to Stretched Clusters; vSAN Stretched Cluster Guide

NEW QUESTION 26

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 29

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 30

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 35

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 39

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 41

An organization plans to implement a new vSAN 8.0 cluster to take advantage of the new features around improved I/O flow, better resiliency, and more efficient disk usage. The vSAN ReadyNodes available for the cluster consist of eight NVMe disks.

How should the organization configure the disk layout?

- A. Use vSAN OSA and create two disk groups with one cache disk and three capacity disks each
- B. Use vSAN ESA and the new Storage pool configuration where all disks contribute to capacity
- C. Use vSAN OSA and the new Storage pool configuration where all disks contribute to capacity
- D. Use vSAN ESA and create two disk groups with one cache disk and three capacity disks each

Answer: B

Explanation:

Using vSAN ESA and the new Storage pool configuration where all disks contribute to capacity is the correct answer because it allows the organization to take advantage of the new features in vSAN 8.0, such as improved I/O flow, better resiliency, and more efficient disk usage. With vSAN ESA, there is no need to create disk groups or designate cache disks, as all disks are treated as capacity disks and use a new algorithm to distribute data across them. This also simplifies the disk management and reduces the overhead of cache management. References:

? VMware vSAN Specialist v2 Exam Preparation Guide, page 6

? What's New in VMware vSAN 8.0

NEW QUESTION 43

When adding a disk to a host that was previously used in a decommissioned vSAN cluster the intended disk does not show among the available devices in disk management.

Which action should be taken prior to assigning the disk on disk management?

- A. Format the existing partition
- B. Create a 1GB metadata partition
- C. Delete all device partitions
- D. Create a VMFS partition

Answer: C

Explanation:

When adding a disk to a host that was previously used in a decommissioned vSAN cluster, the disk may still have some vSAN metadata partitions that prevent it from being recognized by disk management. To resolve this issue, the disk partitions need to be deleted using either ESXCLI or partedUtil commands. This will erase all data on the disk and make it available for use in disk management. References: VMware vSAN Specialist v2 EXAM 5V0-22.23, page 21

NEW QUESTION 45

An administrator is deploying a new two-node vSAN cluster with a shared witness to a remote location.

Which requirement must be met?

- A. The ESXi hosts must have SSDs or NVMe configured for Virtual Flash File System.
- B. The ESXi host's controller cache and advanced features must be disabled.
- C. The ESXi host's drives must be configured in RAID 1 to support Failures to Tolerate of 1.
- D. The ESXi hosts must have a minimum of 64 GBs of memory.

Answer: D

Explanation:

To deploy a new two-node vSAN cluster with a shared witness, the administrator must meet several requirements, one of which is that the ESXi hosts must have

a minimum of 64 GBs of memory. This is because each host must have enough memory to run the VMs and also to support the vSAN metadata overhead. The other options are not requirements for a two-node vSAN cluster with a shared witness. The ESXi hosts do not need SSDs or NVMe for Virtual Flash File System, as they can use any supported storage devices for vSAN. The ESXi host's controller cache and advanced features do not need to be disabled, as they can be used to improve performance and reliability. The ESXi host's drives do not need to be configured in RAID 1, as vSAN uses its own software-defined RAID mechanism to provide Failures to Tolerate.

References: Shared Witness for 2-Node vSAN Deployments; Two-Node Cluster Requirements

NEW QUESTION 47

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 51

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 52

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 54

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 5V0-22.23 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 5V0-22.23 Product From:

<https://www.2passeasy.com/dumps/5V0-22.23/>

Money Back Guarantee

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