

VMware

Exam Questions 5V0-22.23

VMware vSAN Specialist (v2)



NEW QUESTION 1

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 2

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 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 -enable`². After enabling TRIM and UNMAP, the administrator must power off and then power on all VMs that use the vSAN datastore.

NEW QUESTION 4

An administrator has to perform maintenance on one of the hosts in a three-node vSAN Cluster.

Which maintenance mode option will give the administrator the best availability for the VMs with the least effort and data transfer?

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

Answer: D

Explanation:

To perform maintenance on one of the hosts in a three-node vSAN cluster with the best availability for the VMs with the least effort and data transfer, the

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

NEW QUESTION 5

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

NEW QUESTION 6

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

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

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

- A. A new three-node vSAN 8.0 All-Flash Cluster with OSA in each edge location Each application VM configured with a RAID-5 VM storage policy
- B. A new three-node vSAN 8.0 All-Flash Cluster with OSA in each edge location Each application VM configured with a RAID-1 VM storage policy
- C. A new three-node vSAN 8.0 All-Flash Cluster with ESA in each edge location Each application VM configured with a RAID-1 VM storage policy
- D. A new three-node vSAN 8.0 All-Flash Cluster with ESA in each edge location Each application VM configured with a RAID-5 VM storage policy

Answer: B

Explanation:

To satisfy the customer's needs for high-performance, low-latency applications at edge locations, the best deployment option is to use a new three-node vSAN 8.0 All-Flash Cluster with OSA in each edge location and configure each application VM with a RAID-1 VM storage policy. This option will provide the following benefits:

? All-flash clusters offer the highest performance and lowest latency for vSAN, as they use flash devices for both cache and capacity tiers. Flash devices have faster read and write operations than magnetic disks, and they also support advanced features such as deduplication, compression, and encryption.

? OSA stands for One Socket Architecture, which means that each host has only one CPU socket with multiple cores. This reduces the licensing cost and complexity of vSphere and vSAN, as well as the power consumption and cooling requirements of the hosts. OSA also improves the performance of vSAN by eliminating the NUMA effect, which is the latency caused by accessing memory or devices across different CPU sockets.

? RAID-1 is a mirroring technique that creates two copies of each data component and places them on different hosts. This provides high availability and fault tolerance for the application VMs, as they can survive the failure of one host or disk. RAID-1 also offers better performance than RAID-5 or RAID-6, as it does not incur any parity overhead or additional write operations.

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

NEW QUESTION 7

A vSAN administrator is investigating vSAN performance related problems but cannot find any vSAN performance statistics on the cluster summary page.

Why is this situation occurring?

- A. The vRealize Operations Manager is not integrated with vSAN cluster.
- B. The administrator has read-only permissions on the cluster level.
- C. vSAN performance statistics are only available via CLI.
- D. vSAN performance service is not enabled.

Answer: D

Explanation:

The reason why the vSAN administrator cannot find any vSAN performance statistics on the cluster summary page is that the vSAN performance service is not enabled. The vSAN performance service is a feature that collects and analyzes performance metrics and displays them in graphical charts in vCenter. The vSAN performance service must be turned on manually for each vSAN cluster, as it is not enabled by default. The other options are not correct. The integration of vRealize Operations Manager with the vSAN cluster is not required to view vSAN performance statistics, as they are available in vCenter. The administrator's permissions on the cluster level do not affect the visibility of vSAN performance statistics, as they are accessible to any user who can view the cluster. vSAN performance statistics are not only available via CLI, as they can also be viewed in vCenter using the vSAN performance service. References: About the vSAN Performance Service; Enable or Disable the Performance Service

NEW QUESTION 8

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 9

An administrator has successfully deployed a vSAN Stretched Cluster and needs to ensure that any virtual machines that are created are placed in the appropriate site.

Which two steps are needed to complete this task? (Choose two.)

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

Answer: AE

Explanation:

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

NEW QUESTION 10

What is the maximum amount of capacity disks an administrator can have in disk groups on a single vSAN OSA host?

- A. 35
- B. 40
- C. 30
- D. 25

Answer: A

Explanation:

The maximum amount of capacity disks an administrator can have in disk groups on a single vSAN OSA host is 35. This is because a single host can have up to five disk groups, and each disk group can have up to seven capacity disks. Therefore, the maximum number of capacity disks per host is $5 \times 7 = 35$. The other options are not correct, as they are lower than the maximum number of capacity disks per host.

References: Designing and Sizing vSAN Storage; [vSAN ReadyNode Hardware Guidance]

NEW QUESTION 10

An administrator is tasked to create a Kerberos secured NFS v4.1 file share.

Which information is minimally required during the configuration of the File Service?

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

Answer: B

Explanation:

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

NEW QUESTION 14

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 18

The Resyncing Objects view in the vCenter UI reports that some objects are currently resyncing. Which two actions would cause this situation? (Choose two.)

- A. A change to the storage policy is applied to the objects.
- B. DRS is relocating VMs between vSAN nodes.
- C. A host failure occurs in the cluster
- D. HA Virtual Machine Monitoring forced a VM to reboot.
- E. VM snapshot is being deleted.

Answer: AC

Explanation:

Two actions that would cause some objects to be currently resyncing are:

- ? A change to the storage policy is applied to the objects. This action triggers a resynchronization of objects to make them compliant with the new policy settings, such as FTT, RAID level, stripe width, etc. The resynchronization process copies data from one host to another to create or update replicas or parity segments.
- ? A host failure occurs in the cluster. This action causes some objects to become non-compliant with their storage policy, as they lose one or more replicas or parity segments due to the host failure. The resynchronization process rebuilds the missing components on other hosts in the cluster to restore compliance and availability. References: : VMware vSphere Storage Guide, page 129 : Monitor the Resynchronization Tasks in the vSAN Cluster 1 : VMware vSAN Specialist v2 Exam Preparation Guide, page 13

NEW QUESTION 22

A vSAN administrator needs to build a vSAN ESA cluster with RAID-5/FTT 1 adaptive storage policy. What is the absolute minimum number of hosts that need to be part of that vSAN ESA cluster?

- A. 6 hosts
- B. 4 hosts
- C. 5 hosts
- D. 3 hosts

Answer: D

Explanation:

To build a vSAN ESA cluster with RAID-5/FTT 1 adaptive storage policy, the absolute minimum number of hosts that need to be part of that vSAN ESA cluster is 3. This is because the vSAN ESA supports a new RAID-5 erasure coding scheme in a 2+1 configuration, which writes the data in a VM as a stripe consisting of 2 data bits and 1 parity bit, across a minimum of 3 hosts. This scheme can tolerate a single host failure (FTT=1) while consuming 1.5x the capacity of the primary data. This scheme is suitable for smaller vSAN clusters that want to reduce capacity usage without compromising performance¹² References: 1: VMware vSAN Specialist v2 Exam Preparation Guide, page 15 2: Adaptive RAID-5 Erasure Coding with the Express Storage Architecture in vSAN 8 3

NEW QUESTION 23

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 26

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

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

Answer: C

Explanation:

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

NEW QUESTION 31

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 36

An all-flash vSAN ESA cluster contains four nodes. Which two storage policies can the cluster satisfy? (Choose two.)

- A. FTT=3 (RAID-1 Mirroring)
- B. FTT=2 (RAID-1 Mirroring)
- C. FTT=1 (RAID-5 Erasure Coding)
- D. FTT=1 (RAID-1 Mirroring)
- E. FTT=2 (RAID-6 Erasure Coding)

Answer: BE

Explanation:

An all-flash vSAN ESA cluster with four nodes can satisfy the storage policies that require FTT=2 (RAID-1 Mirroring) or FTT=2 (RAID-6 Erasure Coding). These policies mean that the cluster can tolerate two host failures while maintaining data availability and redundancy. RAID-1 Mirroring creates three replicas of each object across different hosts, while RAID-6 Erasure Coding splits each object into four data segments and two parity segments across different hosts. Both policies require at least four hosts in the cluster to meet the FTT=2 requirement. The other options are not correct. An all-flash vSAN ESA cluster with four nodes cannot satisfy the storage policies that require FTT=3 (RAID-1 Mirroring) or FTT=1 (RAID-5 Erasure Coding). These policies mean that the cluster can tolerate three or one host failure respectively, but they require more or less hosts than four to do so. RAID-1 Mirroring with FTT=3 requires at least six hosts in the cluster to create four replicas of each object, while RAID-5 Erasure Coding with FTT=1 requires at least three hosts in the cluster to split each object into two data segments and one parity segment. References: vSAN Express Storage Architecture; RAID Configurations, FTT, and Host Requirements

NEW QUESTION 39

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 `esxcli system maintenanceMode set -e false`.

? 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 40

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 42

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. vSAN 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 47

What is the purpose of the TRIM/UNMAP process?

- A. Collects vSAN log files
- B. Repairs internal cache errors

- C. Deletes orphaned snapshots
- D. Reclaims disk space

Answer: D

Explanation:

The purpose of the TRIM/UNMAP process is to reclaim disk space that is no longer used by the guest operating system or the virtual machine. TRIM and UNMAP are commands that allow the guest operating system to inform the underlying storage layer that certain blocks are no longer in use and can be freed up. This process helps to improve storage efficiency and utilization, especially for thin-provisioned disks that grow dynamically as data is written to them. The other options are not correct. The TRIM/UNMAP process does not collect vSAN log files, repair internal cache errors, or delete orphaned snapshots. These are different tasks that are performed by other tools or processes.

References: Enabling TRIM/UNMAP Commands for VMware Cloud on AWS Clusters; Reclaiming guest OS storage in VMware vSAN 6.7 U1 with TRIM/UNMAP process

NEW QUESTION 52

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 55

Which VMware solution requires vSAN usage?

- A. VMware Cloud Foundation
- B. VMware Horizon
- C. VMware Telco Cloud Automation
- D. VMware Aria Automation

Answer: A

Explanation:

The VMware solution that requires vSAN usage is VMware Cloud Foundation. VMware Cloud Foundation is an integrated software stack that bundles compute virtualization (VMware vSphere), storage virtualization (VMware vSAN), network virtualization (VMware NSX), and cloud management and monitoring (VMware vRealize Suite) into a single platform that can be deployed on premises or as a service within a public cloud. VMware Cloud Foundation relies on vSAN as the primary storage solution for its workload domains, which are logical pools of resources that can be used to run different types of workloads. The other options are not correct. VMware Horizon, VMware Telco Cloud Automation, and VMware Aria Automation are VMware solutions that do not require vSAN usage, although they can benefit from it. VMware Horizon is a platform that delivers virtual desktops and applications across a variety of devices and locations, and it can use any supported storage solution, including vSAN. VMware Telco Cloud Automation is a cloud-native orchestration and automation platform that enables communication service providers to accelerate the deployment and lifecycle management of network functions and services across any network and cloud. It can use any supported storage solution, including vSAN. VMware Aria Automation is not a valid VMware solution name.

References: VMware Cloud Foundation Overview; VMware Horizon Overview; VMware Telco Cloud Automation Overview

NEW QUESTION 59

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 64

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 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 72

An existing vSAN OSA cluster has this specification: Four ESXi hosts with all flash configuration
 Each with two disk groups
 Each disk group with one cache device and four capacity devices There are five more device slots available per host
 The CTO would like to provision new applications, and these will need more capacity and performance.
 Which two methods should be used by the vSAN administrator to meet this goal with the least amount of impact? (Choose two.)

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

Answer: CE

Explanation:

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

other options are incorrect for the following reasons:

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

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

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

? VMware vSAN Specialist v2 Exam Preparation Guide, page 10

? Expanding a vSAN Cluster

NEW QUESTION 74

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 77

After a server power failure, the administrator noticed the scheduled resyncing in the cluster monitor displays objects to be resynchronized under the pending category.

Why are there objects in this category?

- A. The delay timer has not expired.
- B. These objects belong to virtual machines, which are powered off.
- C. Object resynchronization must be started manually.
- D. There are too many objects to be synchronized.

Answer: A

Explanation:

The reason why there are objects in the pending category of the scheduled resyncing in the cluster monitor is that the delay timer has not expired. The delay timer is a configurable setting that determines how long vSAN waits before repairing a non-compliant object after placing a host in a failed state or maintenance mode. The default value is 60 minutes, but it can be changed in the vSAN Services configuration. The pending category displays the objects with the expired delay timer that cannot be resynchronized due to insufficient resources in the current cluster or the vSAN FTT policy set on the cluster not being met. The other options are not correct. These objects do not belong to virtual machines that are powered off, as vSAN resynchronizes all objects regardless of their power state. Object resynchronization does not need to be started manually, as vSAN initiates it automatically when the delay timer expires. There are not too many objects to be synchronized, as vSAN can handle multiple resynchronization tasks in parallel. References: Monitor the Resynchronization Tasks in the vSAN Cluster; About vSAN Cluster Resynchronization

NEW QUESTION 78

.....

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](#)