VMware.5V0-22.23.v2024-01-05.q36

Exam Code:	5V0-22.23
Exam Name:	VMware vSAN Specialist (v2)
Certification Provider:	VMware
Free Question Number:	36
Version:	v2024-01-05
# of views:	461
# of Questions views:	4249
https://www.freecram.net/torrent/VMware.5V0-22.23.v2024-01-05.q36.html	

NEW QUESTION: 1
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 groupmust 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 changingcache devices.
D. The cache disk drives must have a larger capacity.
Answer: (SHOW ANSWER)
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: 2
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: (SHOW ANSWER)
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 problemsfor 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: 3
A vSAN administrator encounters a non-compliant virtual machine and the compliance status of some of its objects is noncompliant. vSAN is able to locate a full replica of 55% of the votes for the noncompliant objects.
What will vSAN do with the virtual machine?
A. Mark the virtual machine as compliant and automatically recover the noncompliant objects
B. Power off the virtual machine
C. Mark the virtual machine as inaccessible as vSAN is not able to locate more than 60% of the votes for the objects
D. Mark the virtual machine as orphaned
Answer: (SHOW ANSWER)
Explanation
If vSAN is able to locate a full replica of 55% of the votes for the noncompliant objects of a virtual machine, vSAN will mark the virtual machine as compliant and automatically recover the noncompliant objects. This is because vSAN uses a quorum-based algorithm to determine object availability and compliance. An object is considered available if more than 50% of its votes are accessible, and compliant if it meets its assigned storage policy requirements. In this case, since 55% of the votes are accessible, vSAN can access a full replica of each object and restore its compliance state by rebuilding any missing or corrupted components. The other options are not correct. vSAN will not power off, mark as inaccessible, or mark as orphaned a virtual machine that has more than 50% of its votes accessible, as these actions would result in unnecessary downtime or data loss.
References: Object States That Indicate Problems in vSAN; Accessibility of Virtual Machines Upon a Failure in vSAN

NEW QUESTION: 4
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 relocatingVMs 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: (SHOW ANSWER)
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: 5
The DevOps team of an organization wants to deploy with persistent storage on a dedicated vSAN cluster. The storage administrator is tasked to configure the vSAN cluster and leverage the vSAN Direct feature.
Which two requirements must the administrator meet to complete this task? (Choose two.)
A. A valid vSAN license for the vSAN cluster
B. HA enabled on the vSAN cluster
C. A dedicated network for vSAN Direct
D. An integration with vSAN File Services
E. Unclaimed disks in the hosts forvSAN Direct
Answer: (SHOW ANSWER)
Explanation
To configure vSAN Direct, the administrator must meet two requirements: a valid vSAN license for the vSAN cluster and unclaimed disks in the hosts for vSAN Direct. A vSAN license is required to enable vSAN features and services, including vSAN Direct. Unclaimed disks are local storage devices that are not used by vSAN or any other service, and can be claimed by vSAN Direct to create datastores for persistent storage. The other options are not requirements for vSAN Direct. HA is an optional feature that can be enabled on any cluster, but is not specific to vSAN Direct. A dedicated network for vSAN Direct is not necessary, as vSAN Direct uses the same network as vSAN. An integration with vSAN File Services is not required, as vSAN Direct does not provide file shares, but block storage. References: Set Up vSAN Direct for vSphere with Tanzu; vSAN Licensing Guide

NEW QUESTION: 6
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. Organizational Unit, User Account. Password

Answer: (SHOW ANSWER)

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: 7

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 theVendor 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 (LEAVE A REPLY)

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: 8
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: (SHOW ANSWER)
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 causeconflicts 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: 9
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: (SHOW ANSWER)
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: 10

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?

Freecram.net

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: (SHOW ANSWER)
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: 11
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: (SHOW ANSWER)
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: 12
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. TanzuKubernetes Grid Integrated Edition
Answer: (SHOW ANSWER)
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: 13
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 thenew 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: (SHOW ANSWER)
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: 14
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. lvSoh ere High Availability
Answer: (SHOW ANSWER)
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: 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 (LEAVE A REPLY)
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 clusters12 References: 1: VMware vSAN Specialist v2 Exam Preparation Guide, page 15 2: VMware vSAN 7 Update 1 - HCI Mesh 3

NEW QUESTION: 16
Refer to the exhibit.
An administrator uses SSH to log into a vSAN ESA host and runs theesxcli vsan debug object overview command.

Object UUID	Group UUID	Version	Size	Used	SPBM Profile	Healthy Components
49413f63-84bd-4aba-2ba6-0050560659c0	1a413f63-a8d1-fafb-0809-0050560659c0	17	0.12 GB	0.01 GB	vSAN Default Storage Policy	7 of 8
1c413f63-4c1a-73bc-9046-0050560659c0	1a413f63-a8d1-fafb-0809-0050560659c0	17	2.00 GB	0.70 GB	vSAN Default Storage Policy	5 of 8
1a413f63-a8d1-fafb-0809-0050560659c0	1a413f63-a8d1-fafb-0809-0050560659c0	17	255.00 GB	0.05 GB	vSAN Default Storage Policy	4 of 8
cf403f63-eec5-da41-8599-0050560659f7	cf403f63-eec5-da41-8599-0050560659f7	17	255.00 GB	0.04 GB	vSAN Default Storage Policy	7 of 8
d0403f63-f7af-45cd-1e8a-0050560659f7	cf403f63-eec5-da41-8599-0050560659f7	17	2.00 GB	0.70 GB	vSAN Default Storage Policy	7 of 8
ef403f63-fe7b-66f0-9d4f-0050560659f7	cf403f63-eec5-da41-8599-0050560659f7	17	0.12 GB	0.01 GB	vSAN Default Storage Policy	5 of 8
db413f63-4ca4-7882-1b50-0050560659f9	db413f63-4ca4-7882-1b50-0050560659f9	17	255.00 GB	0.12 GB	vSAN ESA Default Policy - RAID5	8 of 8
dd413f63-e0e3-929d-9b93-0050560659f9	db413f63-4ca4-7882-1b50-0050560659f9	17	90.00 GB	0.01 GB	vSAN ESA Default Policy - RAID5	5 of 8
e2413f63-4072-62cf-2077-0050560659f9	db413f63-4ca4-7882-1b50-0050560659f9	17	4.00 GB	0.01 GB	vSAN ESA Default Policy - RAID5	8 of 8
f0403f63-e677-850f-db46-0050560659f9	f0403f63-e677-850f-db46-0050560659f9	17	255.00 GB	0.05 GB	vSAN Default Storage Policy	7 of 8
0d413f63-8c58-b213-3866-0050560659f9	f0403f63-e677-850f-db46-0050560659f9	17	0.12 GB	0.01 GB	vSAN Default Storage Policy	4 of 8
f1403f63-365f-559e-8165-0050560659f9	f0403f63-e677-850f-db46-0050560659f9	17	2.00 GB	0.72 GB	vSAN Default Storage Policy	4 of 8
f4403f63-50e3-85c4-ed42-0050560659b4	f4403f63-50e3-85c4-ed42-0050560659b4	17	255.00 GB	3.54 GB	vSAN Default Storage Policy	5 of 9

The administrator notices the Healthy Components column, the last column, is reporting some components are not in a fully healthy state.
What could cause this behavior?
A. New physical disks have been claimed and a rebalance operation is underway.
B. The applied Storage policy has been updated.
C. New VMDKs have been added to multiple VMs, but the storage policy has not finished applying.
D. One host is in maintenance mode with ensure accessibility.
Answer: (SHOW ANSWER)

Explanation
The most likely cause for some components to be not in a fully healthy state is that one host is in maintenance mode with the ensure accessibility option. This option creates temporary durability components on other hosts to maintain the required number of failures to tolerate (FTT) until the original components are restored or rebuilt. These durability components are not considered fully healthy because they do not have full redundancy and might not be compliant with the storage policy. The other options do not explain why some components are not fully healthy, as they do not affect the FTT or the compliance state of the objects.
References: Durability Components; esxcli vsan debug object overview

Valid 5V0-22.23 Dumps shared by ExamDiscuss.com for Helping Passing 5V0-22.23 Exam! ExamDiscuss.com now offer the **newest 5V0-22.23 exam dumps**, the ExamDiscuss.com 5V0-22.23 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com 5V0-22.23 dumps with Test Engine here: https://www.examdiscuss.com/VMware/exam/5V0-22.23/premium/ (**75** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 17
A customer has deployed a new vSAN cluster with the following configuration:
5 x vSAN ReadyNodes
All Flash
12 TB Raw Storage
vSAN 8 is deployed with ESA.
New VMs are configured with a RAID-5 VM policy.
Which statement is accurate?
A. vSAN will use a 2+1 RAID-5 data placement scheme with parity will be used
B. RAID 5 will provide an FTT=2 level of protection in this case
C. vSAN will use a 4+1 RAID-5 data placement scheme with parity will be used
D. vSAN will spread the components across all of the disk groups
Answer: (SHOW ANSWER)
Explanation
vSAN will use a 4+1 RAID-5 data placement scheme with parity will be used is the correct answer because vSAN 8 ESA uses adaptive RAID-5 erasure coding that depends on the number of hosts in the cluster. If the cluster has 6 or more hosts, vSAN will use a 4+1 RAID-5 scheme, where the data is written as a stripe of 4 data bits and 1 parity bit across 5 hosts. This provides a failure tolerance of 1 (FTT=1) and a space efficiency of 1.25x. If the cluster has less than 6 hosts (3 to 5), vSAN will use a 2+1 RAID-5 scheme, where the data is written as a stripe of 2 data bits and 1 parity bit across 3 hosts. This also provides a failure tolerance of 1 (FTT=1) but a space efficiency of 1.5x. In this case, the cluster has 5 hosts, so vSAN will use the 4+1 RAID-5 scheme.

The other options are incorrect for the following reasons:
A, vSAN will use a 2+1 RAID-5 data placement scheme with parity will be used, is incorrect because vSAN will only use this scheme if the cluster has less than 6 hosts but more than 2 hosts. In this case, the cluster has 5 hosts, so vSAN will use the 4+1 RAID-5 scheme.
B, RAID 5 will provide an FTT=2 level of protection in this case, is incorrect because RAID 5 can only provide an FTT=1 level of protection, regardless of the number of hosts or the data placement scheme.
To achieve an FTT=2 level of protection, vSAN would need to use RAID 6 erasure coding, which requires at least 6 hosts in the cluster.
D, vSAN will spread the components across all of the disk groups, is incorrect because vSAN will not necessarily spread the components across all of the disk groups in the cluster. vSAN will only spread the components across as many disk groups as needed to meet the storage policy requirements and to balance the load and capacity. In this case, vSAN will only need to spread the components across 5 disk groups for each stripe of RAID-5 data. References:
VMware vSAN Specialist v2 Exam Preparation Guide, page 11
Adaptive RAID-5 Erasure Coding with the Express Storage Architecture in vSAN 8

NEW QUESTION: 18
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: (SHOW ANSWER)
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. AHost 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: 19
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: (SHOW ANSWER)
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: 20
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?
Freecram.net
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: (SHOW ANSWER)
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 tolerated4 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 orunavailability4 References: 4: VMware vSAN Specialist v2 Exam Preparation Guide, page 13

NEW QUESTION: 21
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: (SHOW ANSWER)
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: 22
Which vSAN maintenance mode option should be used to avoid storage policy non-compliance?
A. Ensure accessibility
B. Partial maintenance mode
C. Full data migration
D. No data migration
Answer: (SHOW ANSWER)
Explanation
To avoid storage policy non-compliance, the vSAN maintenance mode option that should be used is Full data migration. This option evacuates all data from the host to other hosts in the cluster and maintains the current object compliance state. This means that the VM objects will have access to all their replicas and will be compliant with their assigned storage policies. The other

options might result in storage policy non-compliance, as they do not guarantee full data redundancy or policy adherence. Ensure accessibility only migrates the components that are essential for running the VMs, but might not have access to all their replicas.
Partial maintenance mode is not a valid option for vSAN clusters. No data migration does not evacuate any data from the host and might result in VM unavailability or data loss. References: Working with Maintenance Mode; Place a Member of vSAN Cluster in Maintenance Mode

NEW QUESTION: 23
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: C,E (LEAVE A REPLY)
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 afixed 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: 24
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: (SHOW ANSWER)
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: 25
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: (SHOW ANSWER)
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: 26
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: (SHOW ANSWER)
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 met3 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 port3 References: 3: VMware vSAN Specialist v2 Exam Preparation Guide, page 15

NEW QUESTION: 27
A customer wants to validate if Skyline online health is working for vSAN and finds out that Skyline is not fully configured yet.
What two requirements must be met to make sure that Skyline online health will work? (Choose two.)
A. Add the Skyline license into Virtual Center

B. Enable Skyline Health on the vSAN Cluster
C. Enable CEIP and join the program
D. Have a working Internet connection
E. Have vCenter on version 7 or higher
Answer: (SHOW ANSWER)
Explanation
To make sure that Skyline online health will work for vSAN, two requirements must be met: enable CEIP and join the program, and have a working Internet connection. CEIP stands for Customer Experience Improvement Program, which is a voluntary program that collects anonymous product usage data from customers who participate in it. By enabling CEIP and joining the program, customers can benefit from Skyline online health, which provides proactivenotifications and recommendations for software and hardware issues based on VMware Analytics Cloud. A working Internet connection is also required for Skyline online health to communicate with VMware Analytics Cloud and receive online notifications. The other options are not requirements for Skyline online health. References: About the vSAN Skyline Health; Skyline Health

NEW QUESTION: 28
Which VMware solution requires vSAN usage?
A. VMware Cloud Foundation
B. VMware Horizon
C. VMware Telco Cloud Automation
D. VMware Aria Automation
Answer: (SHOW ANSWER)
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: 29
What are two characteristics of the vSAN Data-At-Rest Encryption (DARE)? (Choose two.)
A. it requires Self-Encrypting Drives in order to work.
B. it needs to be enabled together with the vSAN Data-In-Transit encryption.
C. it is Software Defined and works independently of the Cache or Capacity drives installed on the Nodes.
D. it is not supported on Stretched Cluster environments.
E. it continues to operate unaffected during downtime on vCenter Server.
Answer: C,E (LEAVE A REPLY)
Explanation
Two characteristics of the vSAN Data-At-Rest Encryption (DARE) are that it is Software Defined and works independently of the Cache or Capacity drives installed on the Nodes, and that it continues to operate unaffected during downtime on vCenter Server. DARE is a feature that encrypts all data stored on vSAN disks using AES-256 XTS mode. It does not require Self-Encrypting Drives (SEDs) to work, as it uses software-based encryption keys that are generated by an external Key Management Server (KMS) or a vSphere Native Key Provider. DARE also does not depend on the type or size of the disks used in the vSAN cluster, as it encrypts data after all other processing, such as deduplication and compression, is performed.
DARE can function even when the vCenter Server is offline or unavailable, as it uses key persistence to store the encryption keys on the ESXi hosts or in a Trusted Platform Module (TPM). The hosts can access the keys without contacting the KMS or the vCenter Server. The other options are notcorrect, as they do not describe DARE accurately. DARE does not need to be enabled together with the vSAN Data-In-Transit encryption, as they are independent features that can be enabled or disabled separately. Data-In-Transit encryption encrypts data that is transmitted between hosts in a vSAN cluster using secure sockets layer (SSL) certificates. DARE is supported on Stretched Cluster environments, as it can encrypt data across multiple sites using site affinity rules.

NEW QUESTION: 30
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?
Freecram.net
A. Remove the host from ySAN configuration, replace the faulty disks, re-create the storage pool
B. Remove the entire storage pool, install the new devices, re-create the storage pool
C. Remove the host from the cluster, replace the faulty disks, re-add the host to the cluster
D. Remove the devices from the storage pool, replace the storage devices, claim the new devices in vSAN

Answer: (SHOW ANSWER)

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 host4 References: 4: VMware vSphere Storage Guide, page 133 :
VMware vSAN Design and Sizing Guide, page 38

NEW QUESTION: 31

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 confiqured 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: (SHOW ANSWER)

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

Valid 5V0-22.23 Dumps shared by ExamDiscuss.com for Helping Passing 5V0-22.23 Exam! ExamDiscuss.com now offer the **newest 5V0-22.23 exam dumps**, the ExamDiscuss.com 5V0-22.23 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com 5V0-22.23 dumps with Test Engine here: https://www.examdiscuss.com/VMware/exam/5V0-22.23/premium/ (**75** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 32
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: (SHOW ANSWER)
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 ExamPreparation Guide, page 15 2: Adaptive RAID-5 Erasure Coding with the Express Storage Architecture in vSAN 8 3

NEW QUESTION: 33
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: (SHOW ANSWER)
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: 34
A vSAN administrator has two identical VMware vSAN clusters, one for staging workloads and another for production workloads. Due to an unforeseen capacity requirement, the vSAN administrator is tasked with merging the staging vSAN cluster into the production.
Which three actions should the vSAN administrator perform on the staging cluster prior to moving the vSAN nodes to the production cluster? (Choose three.)
A. Disable vSAN Services
B. Delete all Disk Groups
C. Enable File Services
D. Delete all partitions from the capacity disks
E. Mark the disks for partial reservation
F. Remove all capacity drives
Answer: (SHOW ANSWER)
Explanation
The three actions that the vSAN administrator should perform on the staging cluster prior to moving the vSAN nodes to the production cluster are:
Disable vSAN Services: This will stop any vSAN-related operations on the staging cluster, such as resynchronization, rebalancing, or repair. This will also prevent any new virtual machines from being created or migrated to the staging cluster.
Delete all Disk Groups: This will remove all disks from the vSAN cluster and erase all data on them.

This will also free up the disks for use in the production cluster.
Delete all partitions from the capacity disks: This will ensure that there are no remnants of any previous vSAN configuration on the disks. This will also avoid any potential conflicts or errors when adding the disks to the production cluster.
Enabling File Services, marking the disks for partial reservation, and removing all capacity drives are not necessary or recommended actions for this scenario. Enabling File Services would add an unnecessary layer of complexity and overhead to the staging cluster. Marking the disks for partial reservation would reduce the available capacity for vSAN and potentially cause performance issues. Removing all capacity drives would leave only cache disks in the staging cluster, which would not be compatible with vSAN. References:
VMware vSAN Specialist v2 Exam Preparation Guide, page 10

NEW QUESTION: 35
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. Then, remove the disk group with the failed device.
B. Place the disk group into maintenance mode, and select Full Data Migration. Then, physically replace the failed cache device. Afterwards. vSAN will rebuild the disk group automatically.
C Remove the existing vSAN disk group and physically replace thedevice. Thencheck to verify that the ESXi host automatically detects the new device Afterwardsmanually recreate the Disk Group
C. Physically replace the failed cache device, and vSAN will automatically allocate the storage. Then, rebalance the cache layer.
Answer: (SHOW ANSWER)
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: 36
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: (SHOW ANSWER)
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

Valid 5V0-22.23 Dumps shared by ExamDiscuss.com for Helping Passing 5V0-22.23 Exam! ExamDiscuss.com now offer the **newest 5V0-22.23 exam dumps**, the ExamDiscuss.com 5V0-22.23 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com 5V0-22.23 dumps with Test Engine here: https://www.examdiscuss.com/VMware/exam/5V0-22.23/premium/ (**75** Q&As Dumps, **35%OFF Special Discount Code: freecram**)