Pass VMware 5V0-22.23 Actual Free Exam Q&As Updated Dump May 15, 2024 [Q35-Q51

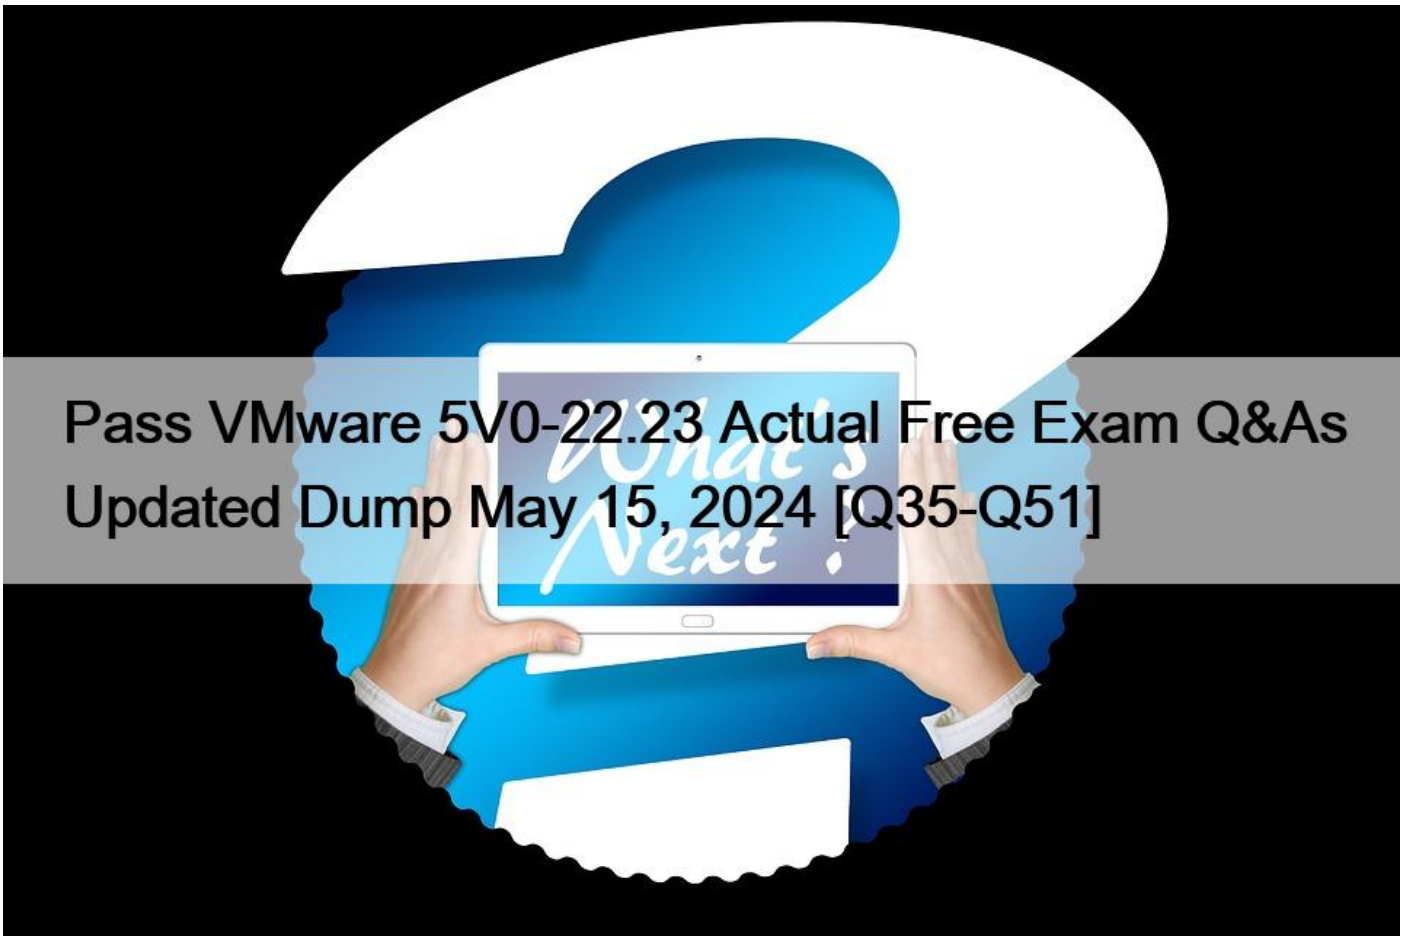

Pass VMware 5V0-22.23 Actual Free Exam Q&As Updated Dump May 15, 2024 Latest 5V0-22.23 Actual Free Exam Updated 75 Questions QUESTION 35

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.)
* Disable vSAN Services
* Delete all Disk Groups
* Enable File Services
* Delete all partitions from the capacity disks
* Mark the disks for partial reservation
* Remove all capacity drives

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

QUESTION 36

What are two prerequisites for using the TRIM and UNMAP capability of vSAN? (Choose two.)
* Deduplication and compression are enabled.
* The vSAN cluster is an all-flash architecture.
* The VM quest operating system supports ATA TRIM or SCSI UNMAP capability
* TRIM and UNMAP is enabled.
* Change the Object Space Reservation to 100.

Explanation

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

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

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

QUESTION 37

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.)
* Replacing all capacity devices with a similar larger device
* Replacing all cache devices with a larger device
* Adding one more disk group per host with the same configuration
* Adding faster cache devices
* Adding an ESXi host with identical device configuration

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

QUESTION 38

Which VMware solution requires vSAN usage?
* VMware Cloud Foundation
* VMware Horizon

* VMware Telco Cloud Automation
* VMware Aria Automation

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

QUESTION 39

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?

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

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)

QUESTION 40

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?
* 6 hosts
* 4 hosts
* 5 hosts
* 3 hosts

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

QUESTION 41

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.)
* Configure disk striping in OSA
* Configure disk striping in ESA
* Enable Force Provisioning in OSA
* Enable deduplication for vSAN
* Create a dedicated Storage Pool in ESA

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 attribute12 References: 1: VMware vSAN Specialist v2 Exam Preparation Guide, page 14 2: VMware vSAN Design and Sizing Guide, page 32

QUESTION 42

How often does the Skyline Health interval validate online if there are new Health Checks available for vSAN?
* Every 1 hour
* Every 4 hours
* Every 24 hours
* Every 12 hours

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

QUESTION 43

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.)
* Create VM/Host groups for the two sites
* Create a single VM/Host group across both sites
* Put the VMs in a vSphere DRS group
* Put the VMs in the correct VM group
* Create a storage policy that includes site affinity rules and apply to VMs

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

QUESTION 44

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 valid vSAN license for the vSAN cluster
* HA enabled on the vSAN cluster
* A dedicated network for vSAN Direct
* An integration with vSAN File Services
* Unclaimed disks in the hosts forvSAN Direct

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

QUESTION 45

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

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

QUESTION 46

Which two actions are recommended when adding a host to a vSAN cluster? (Choose two.)
* Create uniformly-configured hosts
* Disable vSAN performance service
* Disable vSphere Cluster Services
* Disable vSphere High Availability (HA)
* Reference the VMware Compatibility Guide

Explanation

When adding a host to a vSAN cluster, it is recommended to create uniformly-configured hosts and reference the VMware Compatibility Guide. These actions will ensure that the host meets the hardware and software requirements for vSAN, and that it can work seamlessly with the existing hosts in the cluster.

Uniformly-configured hosts have the same number and type of disk groups, cache devices, capacity devices, network adapters, and drivers. The VMware Compatibility Guide provides a list of certified components and firmware versions that are compatible with vSAN. The other options are not recommended, as they can cause disruption or degradation of the vSAN cluster services. Disabling vSAN performance service, vSphere Cluster Services, or vSphere High Availability (HA) can affect the monitoring, availability, and load balancing of the cluster.

QUESTION 47

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?
* vSAN will use a 2+1 RAID-5 data placement scheme with parity will be used
* RAID 5 will provide an FTT=2 level of protection in this case
* vSAN will use a 4+1 RAID-5 data placement scheme with parity will be used
* vSAN will spread the components across all of the disk groups

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

QUESTION 48

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?
* Format the existing partition
* Create a 1GB metadata partition
* Delete all device partitions
* Create a VMFS partition

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

QUESTION 49

In which type of environment is vSAN storage used as a mandatory, primary storage?
* VMware Cloud on AWS
* VMware Horizon
* VMware Aria Automation
* TanzuKubernetes Grid Integrated Edition

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

QUESTION 50

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

Which requirement must be met?
* The ESXi hosts must have SSDs or NVMe configured for Virtual Flash File System.
* The ESXi host’s controller cache and advanced features must be disabled.
* The ESXi host’s drives must be configured in RAID 1 to support Failures to Tolerate of 1.
* The ESXi hosts must have a minimum of 64 GBs of memory.

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

QUESTION 51

What is the minimum required number of hosts to provide data redundancy for a vSAN stretched cluster using dual-site mirroring and local protection with 1 failure – RAID-1 (Mirroring)?
* 3 hosts
* 3 hosts
* 4 hosts
* 6 hosts
Explanation

The minimum required number of hosts to provide data redundancy for a vSAN stretched cluster using dual-site mirroring and local protection with 1 failure – RAID-1 (Mirroring) is six hosts. This is because a vSAN stretched cluster requires at least three hosts per site, and each site must have enough hosts to tolerate one host failure. Therefore, the minimum configuration is three hosts per site, plus one witness host at a third site, for a total of six hosts. References: [VMware vSAN Specialist v2 EXAM 5V0-22.23], page 14

VMware 5V0-22.23 certification exam is designed for IT professionals who want to demonstrate their expertise in vSAN technology. VMware vSAN Specialist (v2) certification exam is intended for those who have a deep understanding of vSAN technology, including its architecture, deployment, and management. VMware vSAN is a software-defined storage solution that enables IT administrators to create highly scalable and resilient virtualized storage infrastructure.

VMware 5V0-22.23 exam is a valuable certification for IT professionals who specialize in VMware vSAN technology. It tests the candidate's ability to deploy, configure, manage, and troubleshoot vSAN environments and covers the latest updates and features of vSAN. By earning the VMware vSAN Specialist (v2) certification, professionals can enhance their career prospects and demonstrate their expertise in vSAN technology.

Online Questions - Valid Practice 5V0-22.23 Exam Dumps Test Questions:
https://www.test4engine.com/5V0-22.23_exam-latest-braindumps.html]