[Q21-Q36 Get Prepared for Your 5V0-22.23 Exam With Actual VMware Study Guide!

Get Prepared for Your 5V0-22.23 Exam With Actual VMware Study Guide! Pass Your Next 5V0-22.23 Certification Exam Easily & Hassle Free

VMware 5V0-22.23 exam consists of 60 multiple-choice questions, which candidates must complete within 105 minutes. 5V0-22.23 exam covers a range of topics, including vSAN architecture, configuration, virtual machine management, storage policies, and performance monitoring. Successful candidates will be awarded the VMware vSAN Specialist certification, which demonstrates their expertise in this area and enhances their career prospects in the field of virtualization and cloud computing.

VMware 5V0-22.23 certification exam is designed for IT professionals who want to validate their skills and knowledge in implementing and managing VMware vSAN solutions. VMware vSAN Specialist (v2) certification exam is the newest version of VMware's vSAN specialist certification and it covers the latest updates and features of vSAN. 5V0-22.23 exam tests the candidate's ability to design, implement, manage, and troubleshoot vSAN solutions in a VMware vSphere environment.

QUESTION 21

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 22

A vSAN administrator is responsible for managing a customer’s production vSAN cluster that is going to be used to provide SMB file shares to a number of host clients. The vSAN administrator must take action so the performance of all services in the production vSAN cluster can be monitored.

Which two services must be enabled for this monitoring to occur? (Choose two.)
* vSAN Performance Diagnostic Service
* iSCSI Target Service
* vSAN File Services
* vSAN Health Service
* vSAN Performance Service

Explanation

To monitor the performance of vSAN File Services, the vSAN administrator must enable both the vSAN File Services and the vSAN Performance Service. The vSAN File Services provides SMB file shares to host clients, while the vSAN Performance Service collects and analyzes performance statistics and displays them in the vSphere Client. The other services are not related to vSAN File Services performance monitoring.

References: VMware vSAN Specialist v2 EXAM 5V0-22.23, page 9, Objective 7.4; [vSAN File Services];

[vSAN Performance Service]

QUESTION 23

A three-node vSAN OSA cluster with business critical intensive I/O workload is running out of capacity. Each host consists of five disk groups with four capacity disks. The administrator needs to expand the capacity of the vSAN datastore as soon as possible.

What should the administrator do?
* Enable Deduplication and Compression on the cluster level

* Add additional capacity by adding a disk on one host and creating a storage pool
* Add additional capacity by addinga vSAN ReadyNode to the cluster
* Add additional capacity disks to each disk group

Explanation

The correct answer is D, add additional capacity disks to each disk group. This is because adding capacity disks to existing disk groups is the fastest and easiest way to expand the capacity of the vSAN datastore without disrupting any ongoing operations or requiring additional hardware. The administrator can add up to five capacity disks per disk group in vSAN OSA, which means each host can have up to 25 capacity disks in total. The administrator should make sure that the new capacity disks are unformatted and not partitioned, so that vSAN can recognize and claim them. The administrator should also manually rebalance the cluster after adding the capacity disks to distribute the data evenly across the new devices. The other options are incorrect for the following reasons:

A, enable Deduplication and Compression on the cluster level, is incorrect because enabling Deduplication and Compression is not a recommended way to expand the capacity of the vSAN datastore. Deduplication and Compression is a space efficiency feature that reduces the logical space consumption of data by eliminating duplicate blocks and applying compression algorithms. However, enabling Deduplication and Compression requires a full data evacuation and resynchronization, which can be disruptive and time-consuming. Deduplication and Compression also introduces additional CPU and memory overhead, which can affect the performance of the cluster. Deduplication and Compression is only supported on all-flash clusters, not on hybrid clusters.

B, add additional capacity by adding a disk on one host and creating a storage pool, is incorrect because creating a storage pool is not supported in vSAN OSA. A storage pool is a new configuration introduced in vSAN 8 ESA, where all disks are treated as capacity disks and use a new algorithm to distribute data acrossthem. This configuration is not compatible with vSAN OSA, which uses a disk group configuration where one disk is designated as a cache disk and the rest are capacity disks. To use a storage pool, the administrator would need to migrate to vSAN 8 ESA on a new cluster with new hardware.

C, add additional capacity by adding a vSAN ReadyNode to the cluster, is incorrect because adding a vSAN ReadyNode to the cluster is not the fastest or easiest way to expand the capacity of the vSAN datastore. A vSAN ReadyNode is a preconfigured server that meets the hardware requirements for running vSAN. Adding a vSAN ReadyNode to the cluster would require additional hardware procurement, installation, and configuration. It would also increase the compute capacity of the cluster, which may not be necessary for the workload. Adding a vSAN ReadyNode would also trigger a resynchronization of data across the cluster, which can affect the performance and availability of the cluster. References:

VMware vSAN Specialist v2 Exam Preparation Guide, page 10

QUESTION 24

An administrator wants to assign a storage policy to a workload on a two-node vSAN OSA cluster consisting of three disk groups each with nested fault domains. The virtual machine must be protected against a disk or disk group failure.

Which two storage policies meet these requirements? (Choose two.)

* RAID-5/FTT 2
* RAID-1/FTT 3
* RAID-6/FTT 2
* RAID-5/FTT 1
* RAID-1/FTT 1

Explanation

To protect a virtual machine against a disk or disk group failure, the storage policy must have a failure tolerance method (FTM) of RAID-1 or RAID-6 and a failure to tolerate (FTT) value of at least 1. RAID-1 mirrors the data across multiple disk groups, while

RAID-6 uses erasure coding to stripe the data and parity information across multiple disk groups. RAID-5 is not suitable for this scenario, as it can only tolerate one disk failure per stripe. FTT 2 or 3 would require more disk groups than available in the cluster.

Therefore, the correct options are C and E. References: 1, page 8; 2, section 3.1

QUESTION 25

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 26

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?
* vSAN File Services
* vSAN HCI Mesh
* vSAN Replication
* vSAN Stretched Clusters

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

QUESTION 27

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

Which scenario would prevent this feature from operating properly?
* Underutilized space is above 25-30% of the total capacity threshold.
* The used space on vSAN datastore exceeds the suggested host rebuild threshold.

* The used space on vSAN datastore exceeds the suggested slack rebuild threshold.
* The physical disk has reached an 80% full reactive rebalance threshold.

Explanation

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

[VMware vSAN Specialist v2 EXAM 5V0-22.23], page 28

QUESTION 28

A customer wishes to host a new range of applications with high-performance requirements, specifically, low latency. The current vSAN platform is based on ReadyNode hardware and uses a vSAN 7.0 U2 hybrid topology configuration.

Which would satisfy the customer’s requirement?

* Deploy the application on a new cluster with vSAN 8.0 ESA using a new hardware design
* Deploy the new applications on the existing cluster with a RAID-6 VM storage policy and an additional stripe width of 4
* Deploy the application on a new cluster with vSAN 8.0 OSA using the existing hybrid configuration
* Perform an in-place upgrade from vSAN 7.0 U2 OSA to vSAN 8.0 ESA

Explanation

Deploying the application on a new cluster with vSAN 8.0 ESA using a new hardware design is the correct answer because it will satisfy the customer’s requirement for low latency. vSAN 8.0 ESA is a new architecture that uses a storage pool configuration where all disks are treated as capacity disks and use a new algorithm to distribute data across them. This improves the I/O flow, reduces the write amplification, and eliminates the cache tier bottleneck. Using a new hardware design with all-flash disks or NVMe disks will further enhance the performance and latency of the application, as these disks have faster read and write speeds than hybrid disks. Deploying the new applications on the existing cluster with a RAID-6 VM storage policy and an additionalstripe width of 4, deploying the application on a new cluster with vSAN 8.0 OSA using the existing hybrid configuration, and performing an in-place upgrade from vSAN 7.0 U2 OSA to vSAN 8.0 ESA are not valid or optimal solutions for this scenario. Deploying the new applications on the existing cluster with a RAID-6 VM storage policy and an additional stripe width of 4 will increase the resiliency and availability of the data, but it will also increase the network traffic, disk space consumption, and parity calculation overhead, which will negatively affect the latency and performance of the application. Deploying the application on a new cluster with vSAN 8.0 OSA using the existing hybrid configuration will not improve the latency significantly, as vSAN 8.0 OSA still uses the same disk group configuration as vSAN 7.0 U2 OSA, where one disk is designated as a cache disk and the rest are capacity disks. The cache disk can still become a bottleneck for high-performance applications, especially if it is not an SSD or NVMe disk. Performing an in-place upgrade from vSAN 7.0 U2 OSA to vSAN 8.0 ESA is not possible, as vSAN ESA requires a different hardware design than vSAN OSA. The existing disk groups need to be deleted and all disks need to be erased before switching to vSAN ESA. References:

[VMware vSAN Specialist v2 Exam Preparation Guide], page 6

What’s New in VMware vSAN 8.0

QUESTION 29

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

* Better Performance

* Faster resynchronization
* Faster snapshot creation
* Better Storage utilization
* Better Availability

Explanation

A durability component is a temporary component that is created when a host or disk group is placed in maintenance mode with the Ensure data accessibility option, or when a host or disk group fails unexpectedly.

A durability component improves the availability of data by maintaining the required number of failures to tolerate (FTT) until the original component is restored or rebuilt. A durability component also speeds up the resynchronization process by reducing the amount of data that needs to be copied. The other characteristics are not applicable to a durability component. References: VMware vSAN Specialist v2 EXAM 5V0-22.23, page

10, Objective 6.8; [Durability Components]

QUESTION 30

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 31

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?

* The disk groupmust be deleted on each physical host in the vSAN OSA cluster to use the NVMe device.
* The disk group does not need to be removed before adding new cache.
* The host must be removed from vSAN OSA cluster before changingcache devices.
* The cache disk drives must have a larger capacity.

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

QUESTION 32

Which vSAN maintenance mode option should be used to avoid storage policy non-compliance?
* Ensure accessibility
* Partial maintenance mode
* Full data migration
* No data migration

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

QUESTION 33

After reviewing various performance charts at a cluster level, an administrator found an individual VM impacting overall performance of the vSAN cluster.

What feature should be used to introspect multiple performance metrics of a single virtual machine?
* esxci
* Skyline Health
* I/O Trip Analyzer
* llOlnsiqht

Explanation

To introspect multiple performance metrics of a single virtual machine, such as latency, throughput, IOPS, and congestion, the

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

QUESTION 34

A vSAN administrator wants to transition from VMware Update Manager to vSphere Lifecycle Manager.

Which element is a mandatory requirement to create an image?
* ESXi Version
* Component
* Firmware and Drivers Add-On
* Vendor Add-On

Explanation

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

QUESTION 35

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 36

The vSphere Client reports that the state of some components stored on the vSAN datastore are in the reconfiguring state.

Which situation causes components to enter this state?
* A host in the cluster enters maintenance mode.
* The cluster is recovering from a vSAN failure.
* The applied storage policy is modified.
* Additional storage capacity is added to the cluster.

Explanation

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

Ace 5V0-22.23 Certification with 75 Actual Questions: https://www.topexamcollection.com/5V0-22.23-vce-collection.html]