[Q33-Q54 Get up-to-date Real Exam Questions for 5V0-22.23 UPDATED [2025

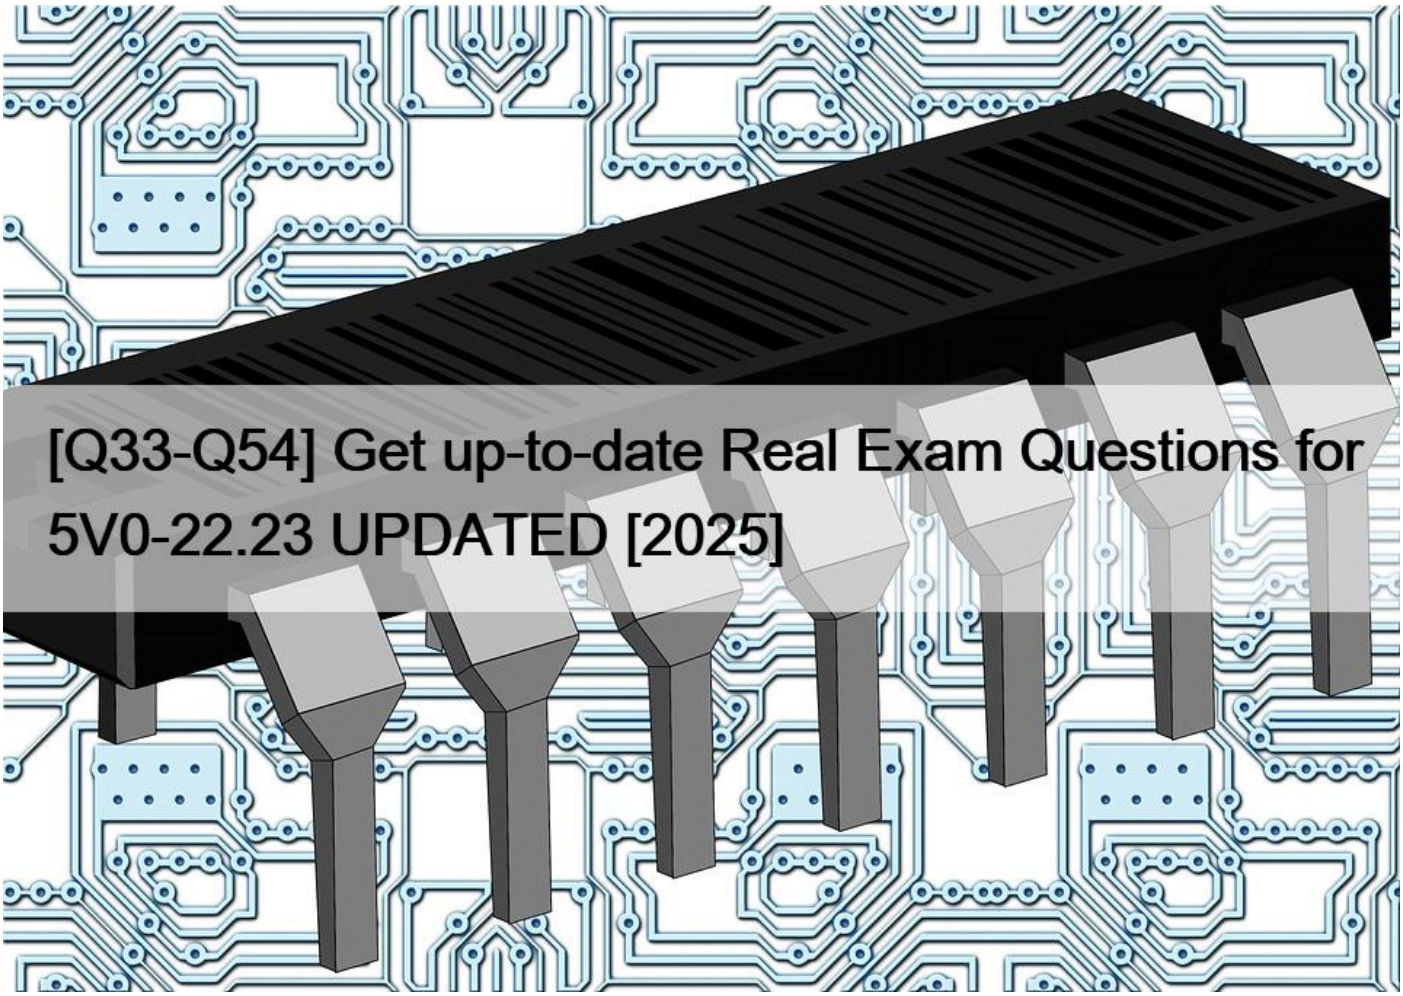

Get up-to-date Real Exam Questions for 5V0-22.23 UPDATED [2025]
Pass VMware 5V0-22.23 Exam in First Attempt Guaranteed

QUESTION 33

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 34

An all-flash vSAN ESA cluster contains four nodes.

Which two storage policies can the cluster satisfy? (Choose two.)
* FTT=3 (RAID-1 Mirroring)
* FTT=2 (RAID-1 Mirroring)
* FTT=I (RAID-5 Erasure Coding)
* FTT=I (RAID-1 Mirroring)
* FTT=2 (RAID-6 Erasure Coding)

Explanation

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

QUESTION 35

What is the purpose of the TRIM/UNMAP process?
* Collects vSAN log files
* Repairs internal cache errors
* Deletes orphaned snapshots
* Reclaims disk space

Explanation

The purpose of the TRIM/UNMAP process is to reclaim disk space that is no longer used by the guest operating system or the virtual machine. TRIM and UNMAP are commands that allow the guest operating system to inform the underlying storage layer that certain blocks are no longer in use and can be freed up. This process helps to improve storage efficiency and utilization,especially for thin-provisioned disks that grow dynamically as data is written to them. The other options are not correct. The TRIM/UNMAP process does not collect vSAN log files, repair internal cache errors, or delete orphaned snapshots. These are different tasks that are performed by other tools or processes. References: Enabling TRIM/UNMAP Commands for VMware Cloud on AWS Clusters; Reclaiming guest OS storage in VMware vSAN 6.7 U1 with TRIM/UNMAP process

QUESTION 36

An administrator is upgrading multiple vSAN Witness nodes with vLCM (single image management) that are used for vSAN Stretched and two-node Clusters.

What two witness node types can the administrator upgrade? (Choose two.)
* Appliance witness node
* Shared witness node
* Nested witness node
* Dedicated witness node
* Physical witness node

Explanation

To upgrade multiple vSAN Witness nodes with vLCM (single image management) that are used for vSAN Stretched and two-node Clusters, the administrator can upgrade two witness node types: appliance witness node and nested witness node. An appliance witness node is a virtual ESXi host that runs on a physical ESXi host and contains the witness components of VM objects stored in the vSAN cluster. A nested witness node is a virtual ESXi host that runs on another virtual ESXi host and contains the witness components of VM objects stored in the vSAN cluster. Both types of witness nodes can be managed by vLCM as independent nodes since vSphere 7.0 Update 3, as long as they are version 7.0 Update 2 or later. The other options are not correct. A shared witness node is a witness node that serves multiple vSAN clusters, which is not supported by vLCM. A dedicated witness node is a witness node that serves only one vSAN cluster, which is not a specific type of witness node. A physical witness node is aphysical ESXi host that contains the witness components of VM objects stored in the vSAN cluster, which cannot be upgraded by vLCM. References: vSphere Lifecycle Manager and the vSAN Witness Hosts; Shared Witness for 2-Node vSAN Deployments

QUESTION 37

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?
* Distribute servers across at least two different racks and configure two fault domains
* Configure disk groups with a minimum of four capacity disks in each server and distribute them across four racks

* Enable deduplication and compression
* Distribute servers across at least three different racks and configure three fault domains

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

QUESTION 38

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 39

An organization wants to implement a virtual desktop infrastructure (VDI) solution on their vSAN storage.

They also need to store their applications running inside the VDI environment on vSAN storage.

Which two end-user computing (EUC) solutions could be implemented to satisfy the requirements of the organization? (Choose two.)

* Agp_ Volumes
* Workspace ONE Access
* Horizon
* Workspace ONE UEM
* Dynamic Environment Manager

Explanation

Horizon and Dynamic Environment Manager are two end-user computing (EUC) solutions that can be implemented on vSAN storage to provide a virtual desktop infrastructure (VDI) solution and store applications running inside the VDI environment. Horizon is a platform that delivers virtual desktops and applications across a variety of devices and locations, while Dynamic Environment Manager is a tool that provides personalization and dynamic policy configuration across any virtual, physical, and cloud-based Windows desktop environment. The other solutions are not directly related to VDI or application storage on vSAN.

References: VMware vSAN Specialist v2 EXAM 5V0-22.23, page 8, Objective 3.5; [Horizon]; [Dynamic Environment Manager]

QUESTION 40

A vSAN administrator is using the vSAN ReadyNode Sizer to build a new environment. While entering the cluster configurations, a fellow colleague inquires about the Operations Reserve option.

What is the purpose of using this option?
* Provides space for internal operations
* Configures space for external operations
* Reserves space for tolerating failures
* Allocates space forvSAN uparades

Explanation

The purpose of using the Operations Reserve option in the vSAN ReadyNode Sizer is to provide space for internal operations such as deduplication, compression, encryption, snapshots, clones, and rebalancing. The Operations Reserve is calculated as a percentage of the total usable capacity of the vSAN cluster. The default value is 30%, but it can be adjusted based on the expected workload characteristics and data services requirements. The other options are not correct, as they do not describe the Operations Reserve option. Configuring space for external operations, reserving space for tolerating failures, and allocating space for vSAN upgrades are not part of the Operations Reserve option. References: 2, section 2; , section 3

QUESTION 41

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 42

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 43

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?
* Preferred and secondary sites
* vCenter High Availability
* Distributed Resource Scheduler
* IvSoh ere High Availability
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

QUESTION 44

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?
* Physically replace the failed cache device, and vSAN will automatically create a new disk group. Then, remove the disk group with the failed device.
* 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
* Physically replace the failed cache device, and vSAN will automatically allocate the storage. Then, rebalance the cache layer.
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

QUESTION 45

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 46

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 47

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?
* No changes occur until the policy is reapplied.
* The changes are queued for 60 minutes.
* The policy change is rejected immediately.
* The updated policy is serially applied to the virtual machines.

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

QUESTION 48

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 49

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 50

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.)

* Add the Skyline license into Virtual Center
* Enable Skyline Health on the vSAN Cluster
* Enable CEIP and join the program
* Have a working Internet connection
* Have vCenter on version 7 or higher

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

QUESTION 51

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

Why are there objects in this category?

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

Explanation

The reason why there are objects in the pending category of the scheduled resyncing in the cluster monitor is that the delay timer has not expired. The delay timer is a configurable setting that determines how long vSAN waits before repairing a non-compliant object after placing a host in a failed state or maintenance mode. The default value is 60 minutes, but it can be changed in the vSAN Services configuration. The pending category displays the objects with the expired delay timer that cannot be resynchronized due to insufficient resources in the current cluster or the vSAN FTT policy set on the cluster not being met. The other options are not correct.

These objects do not belong to virtual machines that are powered off, as vSAN resynchronizes all objects regardless of their power state. Object resynchronization does not need to be started manually, as vSAN initiates it automatically when the delay timer expires. There are not too many objects to be synchronized, as vSAN can handle multiple resynchronization tasks in parallel.
References: Monitor the Resynchronization Tasks in the vSAN Cluster; About vSAN Cluster Resynchronization

QUESTION 52

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 53

What is the purpose of host rebuild reserve in vSAN?
* Reserves space for internal operations
* Reserves space in case of single host failure
* Stores vSphere HA heartbeats
* Allocates capacity for vCLS

Explanation

The host rebuild reserve is a feature that allows vSAN to reserve space in the cluster for vSAN to be able to repair in case of a single host failure. This reservation is set to one host worth of capacity, which means that if one host in the vSAN cluster fails and no longer contributes storage, there is still sufficient capacity remaining in the cluster to rebuild and re-protect all vSAN objects. This feature prevents the creation of new VMs or powering on VMs if such operations consume the reserved space. By default, the host rebuild reserve is disabled, but it can be enabled in the vSAN Services configuration. The other options are not related to the hostrebuild reserve. References: vSAN Capacity Management in v7.0U1; Configure Reserved Capacity

QUESTION 54

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 change to the storage policy is applied to the objects.
* DRS is relocatingVMs between vSAN nodes.
* A host failure occurs in the cluster
* HA Virtual Machine Monitoring forced a VM to reboot.
* VM snapshot is being deleted.

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