Vendor:VMware

Exam Code:2VB-601

Exam Name:VMware Specialist: vSAN 6.x Exam

Version: Demo

QUESTION 1

A cache tier device experiences a permanent disk failure.

Which of the below represents a valid process for replacing the failed disk?

A. 1. The hot spare disk replaces the failed disk immediately.

2. Physically replace the failed disk.

B. 1. Using the vSphere Web Client, replace the failed disk with the hot spare disk.

2. Physically replace the failed disk.

C. 1. Using the vSphere Web Client, delete the failed disk from the vSAN cluster.

2.

Physically replace the failed disk.

3.

Using the vSphere Web Client, re-create the disk group.

D. 1. Physically replace the failed disk and vSAN automatically takes care of the rest of the process.

Correct Answer: C

When a flash device failure occurs, before physically removing the device from a host, you must decommission the device from Virtual SAN. The decommission process performs a number of operations in order to discard disk group memberships, deletes partitions and remove stale data from all disks.

How to do it. Flash Device Decommission Procedure from the vSphere Web Client Log on to the vSphere Web Client Navigate to the Hosts and Clusters view and select the cluster object Go to the manage tab and select Disk management under the Virtual SAN section Select the disk group with the failed flash device Select the failed flash device and click the delete button

References:https://blogs.vmware.com/storage/2014/12/02/vmware-virtual-san-operations-replacing-diskdevices/

QUESTION 2

An administrator is designing a vSAN cluster for a number of new workloads. The design calls for the use of erasure coding with RAID 5.

Which statement is true about what is needed to meet the needs of this scenario?

A. vSAN Standard licensing is sufficient

B. vSAN Enterprise licensing is required

C. An all-flash configuration is required.

D. Six nodes are required at a minimum.

Correct Answer: C

RAID 5 or RAID 6 erasure coding is available only on all-flash disk groups.

References: https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID-6D8185558DE8-4F06-9498-66903FB9C775.html

---

QUESTION 3

For which scenario is network multicast required?

A. Clusters with IGMP disabled on all virtual switches.

B. Network environments where IGMP is disabled.

C. Clusters that have more than four fault domains configured.

D. Clusters running versions of vSAN prior to version 6.6.

Correct Answer: D

Explanation: vSAN 6.6 simplifies design and deployment by removing the need for multicast network traffic (required for versions of vSAN prior to 6.6). When upgrading from a previous version of vSAN to vSAN 6.6, multicast is required until all hosts

---

QUESTION 4

Which two requirements must be satisfied for a v object protected by a RAID-1 mirroring policy to be accessible? (Choose two.)

A. Primary level of failures to tolerate is set to 1 or higher.

B. More than 50% of the object\\'s votes must be available.

C. Every component must have at least two votes.

D. At least one full replica of the object must be available.

E. vSphere HA must be enabled.

Correct Answer: BD

QUESTION 5

Consider the following vSAN host configuration:

1.

Each host contains one vSAN disk group

2.

All drives in the vSAN disk group are attached to the same storage controller

3.

All virtual machines are assigned the Virtual SAN Default Storage Policy, which has not been modified

4.

ESXi is installed on and running from a drive connected connected to a separate storage controller

5.

vSphere HA is enabled

What happens when the storage controller with the vSAN drives attached fails?

A. All components on the drives affected by the storage controller failure are marked "Offline". vSphere HA restarts all virtual machines running on the host with the storage controller failure.

B. vSphere HA restarts all virtual machines running on the host with the storage controller failure. vSAN components affected by the storage controller failure are marked "Repairing" until the virtual machines are back online.

C. All components on the capacity affected by the storage controller failure are marked "Stale". vSAN waits 60 minutes before attempting to rebuild the affected components on other healthy hosts in the cluster.

D. All components on the capacity affected by the storage controller failure are marked "Degraded". vSAN attempts to rebuild all components affected by the failure on the other healthy hosts in the cluster.

Correct Answer: D

QUESTION 6

Which statement is true about two-node vSAN cluster configuration?

A. Two-node configurations do not require a witness host appliance.

B. Only the \\'Ensure data accessibility\\' option can be used when entering maintenance mode.

C. Two witness host appliances can be simultaneously used for redundancy.

D. When a host fails, vSAN rebuilds data on the other host to protect against another failure.

Correct Answer: B

QUESTION 7

How does vSAN respond when a vSAN drive is pulled from a server chassis or a host if a vSAN cluster fails without warning?

A. Components on the affected device(s) are marked as absent. vSAN will attempt to rebuild the components after the VSAN.ClomRepairDelay timer, which is set to 60 minutes by default, expires.

B. Components on the affected device(s) are marked as inaccessible. vSAN will attempt to rebuild the components the

next time a host enters maintenance mode.

C. Components on the affected device(s) are marked as degraded. vSAN will attempt to rebuild the components the next time a host enters maintenance mode.

D. Components on the affected device(s) are marked as healthy. vSAN will attempt to rebuild the components after the VSAN.ComponentRecovery timer, which is set to 60 minutes by default, expires.

E. Components on the affected device(s) are marked as stale. vSAN will attempt to rebuild the components after the VSAN.ComponentRecovery timer, which is set to 90 minutes by default, expires.

Correct Answer: A

QUESTION 8

What are two main advantages of using multiple disk groups within each host? (Choose two.)

A. Performance

B. Backward compatibility

C. Cost

D. Redundancy

Correct Answer: AD

http://www.yellow-bricks.com/2014/05/22/one-versus-multiple-vsan-diskgroups-per-host/

QUESTION 9

Which are two characteristics of a RAID-1 storage configuration? (Choose two.)

A. Full copies of data are duplicated across devices for redundancy.

B. Performance is slightly reduced due to in-memory write amplification.

C. Parity data is used to achieve higher performance.

D. More raw storage device capacity is required than a RAID-5 configuration to achieve the same levels of availability.

Correct Answer: AB

QUESTION 10

Which three of the listed performance metrics can be observed for vSAN backend traffic? (Choose three.)

A. Latency

B. Throughput

C. Delayed normalized IOPS

D. Packets per second

E. IOPS

Correct Answer: ABE

QUESTION 11

A virtual machine named VM01 is running on Host01 on a four-node vSAN cluster. vSphere HA is enabled for the cluster. vSphere HA host isolation response is set to Power off and restart VMs. VM01 has a storage policy assigned to it that contains the rule Primary level of failures to tolerate = 1.

A network failure occurs that causes a vSAN network partition. Host01 is isolated from the rest of the cluster. The other three nodes are NOT affected by the network failure.

How does this network failure impact VM01?

A. vSphere HA attempts to restart VM01 on Host01 after the number of minutes configured in the VSAN.ClomRepairDelay advanced setting has elapsed.

B. vSphere HA immediately attempts to restart VM01 on Host01.

C. vSAN rebuilds all components that belong to VM01 on other nodes in the cluster. After the rebuild is complete vSphere HA attempts to restart VM01 on a node other than Host01.

D. vSphere HA powers off VM01 and attempts to restart it on a node other than Host01.

Correct Answer: C

vSAN will wait for 60 minutes by default and then rebuild the affected data on other hosts in the cluster. The 60-minute timer is in place to avoid unnecessary movement of large amounts of data. As an example, a reboot takes the host offline for approximately 10 minutes. It would be inefficient and resource intensive to begin rebuilding several gigabytes or terabytes of data when the host is offline briefly. vSphere HA is tightly integrated with vSAN. The VMs that were running on a failed host are rebooted on other healthy hosts in the cluster in a matter of minutes. A click-through demonstration of this scenario is available here: vSphere HA and vSAN 50 VM Recovery.

Recommendation: Enable vSphere HA for a vSAN cluster. Reference: vSAN Frequently Asked Questions (FAQ).pdf (Page 11)

QUESTION 12

Consider the following vSAN stretched cluster scenario:

1.

Site A is the preferred site, Site B is the secondary site

2.

Site C is running the vSANwitness host virtual appliance

3.

A virtual machine named VM01 is located on the vSAN datastore

4.

VM01 is running on a host at SiteB

5.

The vSAN default storage policy is assigned to VM01

6.

The vSAN default storage policy has NOT been modified

7.

All aspects of the cluster are functioning properly

Where are reads and writes for VM01 performed?

A. Reads are performed at Site B, writes are performed synchronously at Site A and Site B.

B. Reads and writes are performed at Site A and Site B using a round-robin algorithm.

C. Reads and writes are performed at Site A since it is the preferred site. Changes to Site A are replicated asynchronously to Site B

D. Reads are performed at Site B, writes are performed at Site B. Changes to Site B are replicated asynchronously to Site A through the vSAN witness host.

Correct Answer: C