

2VB-601^{Q&As}

VMware Specialist: vSAN 6.x Exam

Pass VMware 2VB-601 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/2vb-601.html>

100% Passing Guarantee
100% Money Back Assurance

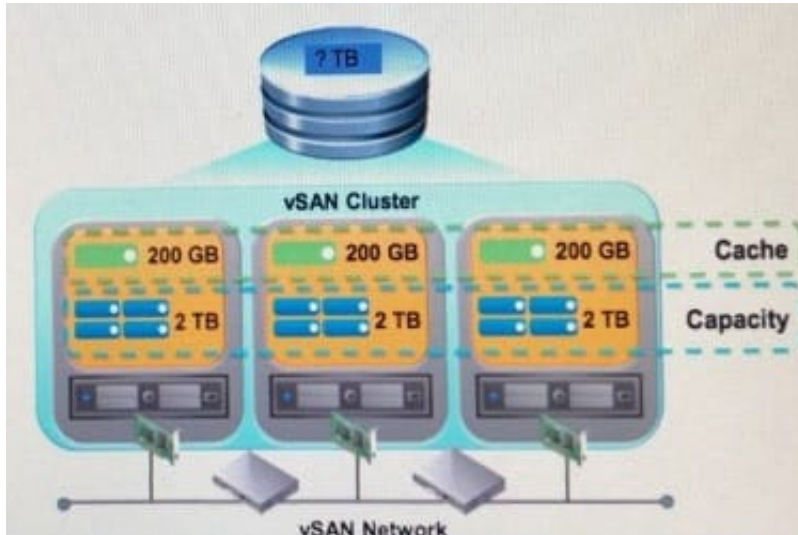
Following Questions and Answers are all new published by VMware
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

Exhibit:



View the exhibit.

The following are configuration details for a three-node all-flash vSAN cluster

1.

Each node is identical in number of drives and disk controllers

2.

Each node has 4 x 500 GB SSD Capacity Tier, 1 x 200 GB SSD cache tier

What is the raw capacity of this cluster as configured?

A. 6.6 TB

B. 3 TB

C. 6 TB

D. 4 TB

Correct Answer: C

To determine the raw capacity of a Virtual SAN datastore, multiply the total number of disk groups in the cluster by the size of the capacity devices in those disk groups, and subtract the overhead required by the Virtual SAN on-disk format.

Reference: <https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID581D2D5C-A88F-4318-A8B3-5A5F343F1247.html>

QUESTION 2

What is the purpose of the cache tier in an all-flash vSAN configuration?

- A. The cache tier is used only for write buffering.
- B. The cache tier is used only for read cache.
- C. 70% of the cache tier is used for read cache and 30% is used for write buffering.
- D. 50% of the cache tier is used for read cache and 50% is used for write buffering.

Correct Answer: A

Explanation: all-flash vSAN configurations continue to use the cache tier as a write buffer References: <https://www.vsan-essentials.com/chapter-5-architectural-details>

QUESTION 3

When using vSAN in a two-node direct connect configuration, how do data nodes communicate with the vSAN witness appliance?

- A. The vCenter Server acts as a proxy between data nodes and the vSAN witness.
- B. An alternate VMkernel interface that has connectivity to the vSAN witness must be tagged with a vSAN traffic type of "witness" VM
- C. Two-node direct connect does NOT require connectivity to the vSAN witness appliance.
- D. If the data node management VMkernel interfaces that are tagged for management traffic only have connectivity with the vSAN witness, they will provide communication with the vSAN witness.

Correct Answer: B

Traditional vSAN 2 Node configurations require connectivity between the vSAN tagged VMkernel ports, and the vSAN Witness Appliance's vSAN tagged VMkernel port. 2 Node Direct Connect Witness Traffic Separation provides the ability to directly connect vSAN data nodes in a 2 Node configuration. Traffic destined for the Witness host can be tagged on an alternative interface from the directly connected vSAN tagged interface.

QUESTION 4

What is the minimum number of components that comprise a 100GB vSAN object with a RAID-5 erasure coding policy assigned to it?

- A. 3
- B. 9
- C. 6
- D. 4

Correct Answer: D

QUESTION 5

Which three of the listed items are required when configuring a vSAN cluster? (Choose three.)

- A. Each host in the vSAN cluster has a VMkernel network adapter for vSAN traffic.
- B. All hosts contributing storage to the cluster have at least one cache and one capacity device.
- C. A valid license for vSAN is applied.
- D. RDMA over Converged Ethernet (RoCE) is enabled on the network switches that connect vSAN hosts.
- E. Storage controllers are configured as RAID-0 with 100% write cache.

Correct Answer: CDE

QUESTION 6

Which vSAN cluster design should be consider for maximum availability?

- A. three-node configuration with multiple disk groups
- B. two-node configuration with a single disk groups
- C. four or more-node configuration with multiple disk groups
- D. two-node configuration with multiple disk groups

Correct Answer: C

QUESTION 7

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 8

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 9

With the Default vSAN Storage Policy applied to all virtual machines, which statement is true about a vSAN three-node configuration?

- A. RAID-5/6 erasure coding can be enabled to save space in an all-flash vSAN configuration.
- B. It can migrate all data from a node during maintenance.
- C. It requires a witness appliance.
- D. It can tolerate only one host failure.

Correct Answer: A

QUESTION 10

The following are configuration details for a three-node vSAN cluster:

- 1.

Each host has two disk groups and each disk group has 7x1.2TB 10K SAS drives.

2.

Each host has 16.8TB of raw capacity and the three-node cluster has 50.4TB of raw capacity.

3.

Deduplication and compression are not enabled.

Based on the following FTT Values in the storage policy, what is the total usable capacity?

Storage Policy	Usable Capacity
FTT = 0	?
FTT = 1	?
FTT = 2	?
FTT = 3	?

A.

Storage Policy	Usable Capacity
FTT = 0	12.6TB
FTT = 1	16.8TB
FTT = 2	25.2TB
FTT = 3	50.4TB

B. Cannot be calculated until vSAN datastore is online.

C.

Storage Policy	Usable Capacity
FTT = 0	50.4TB
FTT = 1	50.4TB
FTT = 2	50.4TB
FTT = 3	50.4TB

D.

Storage Policy	Usable Capacity
FTT = 0	50.4TB
FTT = 1	25.2TB
FTT = 2	16.8TB
FTT = 3	12.6TB

A. Option A

B. Option B

C. Option C

D. Option D

Correct Answer: A

QUESTION 11

Which storage policy is automatically created by Horizon View?

- A. Level of failures to tolerate
- B. Checksum
- C. Horizon Default Storage Policy
- D. Linked Clone

Correct Answer: D

QUESTION 12

In order to support Oracle Real Application Cluster (RAC), which change must be made to each applicable virtual machine?

- A. Set virtual disk sharing to multiwrite mode.
- B. Enable vApp options.
- C. Change SCSI controller type from virtual to physical.
- D. Enable virtual machine Memory Hot Plug.

Correct Answer: A

QUESTION 13

If the Primary level of failures to tolerate policy is changed from 2 to 3 without shutting down the virtual machine, what will happen when the policy is applied?

- A. It will fail with an alert. Primary level of failures to tolerate=3 is not supported with mirroring.
- B. It will fail with an alert. The policy of a running VM cannot be dynamically changed.
- C. vSAN will try to create an additional mirror of the VM's disk components, as long as there is a sufficient number of fault domains and available capacity.
- D. vSAN will switch its failure to tolerate method for any VMs with that policy, to optimize for space. Primary level of failures to tolerate=3 can only be accomplished with erasure coding.

Correct Answer: C

Incorrect Answers:

D: RAID 5 or RAID 6 erasure coding does not support a Primary level of failures to tolerate value of 3.

Note: Primary level of failures to tolerate Defines the number of host and device failures that a virtual machine object can tolerate. For n failures tolerated, each piece of data written is stored in n+1 places, including parity copies if using RAID 5 or RAID

6.

When provisioning a virtual machine, if you do not choose a storage policy, vSAN assigns this policy as the default virtual machine storage policy.

If fault domains are configured, 2n+1 fault domains with hosts contributing capacity are required. A host, which is not part of any fault domain is considered its own single-host fault domain.

Default value is 1. Maximum value is 3.

References: <https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID-08911FD32462-4C1C-AE81-0D4DBC8F7990.html>

QUESTION 14

Which two statements are true? (Choose two.)

- A. Storage policies define VM storage requirements, such as performance and availability
- B. vSAN will not automatically assign a storage policy to a VM
- C. A default storage policy must be created for vSAN by the administrator
- D. The default vSAN storage policy is used unless a different storage policy is selected
- E. The administrator uses the ESX host console to view, create, and modify policies

Correct Answer: AD

QUESTION 15

A 12-node vSAN cluster has been deployed across three server racks in a data center.

1.

There are four nodes in each server rack.

2.

A vSAN fault domain has been configured for each server rack for a total of three fault domains.

3.

All virtual machines in the cluster have the vSAN default storage policy assigned to them.

4.

This policy has not been modified.

5.

vSphere HA is enabled for the cluster.

6.

One of the server racks loses power.

How does this affect virtual machines running on this cluster?

A. Since there are four nodes in each server rack, erasure coding is used to rebuild data that was lost due to the server rack power failure. vSphere HA monitors the rebuild process and powers on all affected virtual machines after the rebuild process is complete.

B. Virtual machine data in the failed server rack is rebuilt immediately using parity information. vSphere HA powers off and restarts all of the virtual machines to rebalance storage utilization across the cluster.

C. All virtual machine data remains accessible. Virtual machines running on nodes in the failed server rack are restarted on nodes in the other two server racks. Virtual machines in the other two server racks continue to run.

D. Virtual machines running on nodes in the failed server rack remain offline for 60 minutes. Virtual machines in the other two server racks continue to run and are marked "Degraded".

Correct Answer: C

[2VB-601 VCE Dumps](#)

[2VB-601 Study Guide](#)

[2VB-601 Braindumps](#)