



- Expert Verified, Online, **Free**.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

Which two vSAN Health Check items are included for vSphere Update Manager integration? (Choose two.)

- A. vSAN build recommendation
- B. vSAN object health
- C. Online health connectivity
- D. vSAN release catalog up-to-date
- E. Performance data collection

Suggested Answer: AC

Reference:

<https://storagehub.vmware.com/t/vmware-vsan/vsan-6-7-update-1-technical-overview/vsphere-update-manager-integration-7/>

Community vote distribution

AD (100%)

 **adelbelkis2** Highly Voted 5 years, 1 month ago

A & D

<https://storagehub.vmware.com/t/vsan-6-7-update-1-technical-overview/vsphere-update-manager-integration-7/>

upvoted 8 times

 **Lazylinux** Most Recent 3 years, 4 months ago

Selected Answer: AD

Is correct

upvoted 1 times

 **Virtual_ITTech** 3 years, 5 months ago

A and D is the correct answer

upvoted 1 times

 **Lazylinux** 3 years, 5 months ago

A&D for sure

upvoted 1 times

 **ZakirK** 3 years, 5 months ago

Selected Answer: AD


As per the VMware documentation

upvoted 1 times

 **Char250** 3 years, 8 months ago

based on this KB <https://kb.vmware.com/s/article/60219> the answers are A & D

upvoted 1 times

 **Labeeb** 3 years, 10 months ago

Correct answers based on the discussions

1 A D

2 A

3 D

4 A

5 C

6 B

7 A

8 A C

9 A C

10 A E

11 A

12 B C

13 A C

14 D
15 A
16 A C
17 C
18 C
19 C
20 C D
21 A E
22 B
23 C
24 C D
25 B
26 C
27 A
28 C D
29 B E
30 A
31 D
32 B
33 C
34 A
35 B
36 B
37 D
38 D
39 D
40 A
41 B E
42 D
43 A
44 A
45 D
46 D
47 D
48 C D
49 D
50 A D
51 C
52 A
53 B
54 A
55 C D
56 B
57 D
58 C
59 C
60 A
61 D
62 B
63 D
64 C
65 B
66 A
67 D
68 A
69 A B
70 B
71 C D

72 C
73 C
74 C
75 A
76 C E
77 A
78 A
79 A
80 C
81 A
82 D
83 B
84 B D
85 D
86 C
87 D
88 C
89 A
90 A
91 D
92 A
93 C
94 B
95 C
96 D
97 C
98 A
99 A C
100 B

upvoted 3 times

  **kskalvik** 3 years, 9 months ago

Based on this i passed with 460

upvoted 1 times

  **MarcusMaia** 4 years, 4 months ago

A & C are correct ... and here is why:

D is under A on vSAN Health Check so once you answer A you are already covering D.



Then we have to choose among B, C and E. (Does anyone go for B or E? lol)

Looking into this KB (<https://kb.vmware.com/s/article/60219>), you may see that inside Hardware Compatibility you have the "vSAN HCL DB up-to-date" and there are 2 ways for that: Manually upload the DB File or keeping the Online Health 'Connectivity' enabled.

And that's the point ... The question doesn't say that VUM requires the results of Online Health Check but that you shall have the connectivity working properly in order to have the "vSAN HCL DB up-to-date"

I hope this helps to clarify.

upvoted 1 times

  **ms200** 4 years, 5 months ago

the questions that are there are reliable for the exam?

Thanks in advance

upvoted 3 times

  **alanp** 4 years, 5 months ago

A and D

<https://cloud.ibm.com/docs/vmwaresolutions?topic=vmwaresolutions-vum-updating-vSAN>

upvoted 1 times

🗨️ 👤 **guruto** 4 years, 5 months ago

A & C

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vsan-monitoring.doc/GUID-B0A8BF17-E3FB-421A-AC1A-8C1EC27294D5.html>

upvoted 1 times

🗨️ 👤 **LookingUp** 4 years, 7 months ago

A & D

<https://core.vmware.com/resource/vsan-operations-guide#sec13-sub1>

Health checks

The vSAN health service provides a number of health checks to ensure the consistency of the hypervisor across the cluster. This is an additional way to alert for potential inconsistencies. The vSAN health checks may also show alerts of storage controller firmware and driver inconsistencies.

Notable health checks relating to updates include:

Customer Experience Improvement Program (necessary to send critical updates regarding drivers and firmware)

vCenter server up-to-date

vSAN build recommendation engine health

vSAN build recommendation

vSAN release catalog

vCenter state is authoritative

vSAN software version compatibility

vSAN HCL DB up-to-date

SCSI controller is VMware certified

Controller is VMware certified for ESXi release

Controller driver is VMware certified

Controller firmware is VMware certified

vSAN firmware provider health

vSAN firmware version recommendation

upvoted 2 times

🗨️ 👤 **jasonv** 4 years, 10 months ago

agree , A and C vsan health and build

upvoted 2 times

🗨️ 👤 **alfred** 4 years, 11 months ago

I think A and C is correct..the release catalog was under on vSan build recommedndation

upvoted 2 times

🗨️ 👤 **Mzoear** 5 years ago

answer is A & D

upvoted 3 times



What is a benefit of using vSphere Distributed Switches in a vSAN network?

- A. Simplifies network abstraction
- B. Enables network redundancy
- C. Supports faster network adapters
- D. Provides the ability to adjust the MTU

Suggested Answer: D

Community vote distribution

A (100%)

  **Mimmo2020** Highly Voted 4 years, 11 months ago

Information Got from the VMware deploy manual training :

• Simplifies management by abstracting networks. • Reduces the amount of manual work and opportunity for mistakes. • Allows the use of Network I/O Control. • vSAN includes distributed switch licensing.

I guess the correct answer is A

upvoted 13 times

  **Darius_Th3D0G** Highly Voted 5 years, 1 month ago

I'm not sure if D is the correct answer. Standard Switches also allow configuring MTU.


Maybe A is the correct answer.

upvoted 7 times

  **Lazylinux** Most Recent 3 years, 5 months ago

I agree A is the answer and it is NOT vSAN specific but more the features of Distributed switch in all circumstances

upvoted 1 times

  **ZakirK** 3 years, 5 months ago

Selected Answer: A



Rest all options are invalid.

upvoted 1 times

  **atti9** 3 years, 10 months ago

B, C and D is also covered by standard vswitches. So A.

upvoted 1 times

  **diegof1** 4 years, 4 months ago

A is the correct answer.

vSphere Distributed Switch has several benefits over standard switches:

- Simplifies management by abstracting networks.
- Reduces the amount of manual work and opportunity for mistakes.
- Allows the use of Network I/O Control.
- vSAN includes distributed switch licensing.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - vSAN with vSphere Distributed Switch section

upvoted 3 times

  **jasonv** 4 years, 10 months ago

i'll go with A

upvoted 5 times

  **Mzoear** 5 years ago

i think it's A

upvoted 6 times

  **mariogg85** 5 years, 1 month ago

If I had to choose an answer would be the answer A, but by discard. The answer should be "The benefit of using vSphere Distributed Switch is that it supports Network I/O Control which supports prioritization of bandwidth if contention occurs."

upvoted 4 times

A single disk in a vSAN disk group suffers from an unrecoverable hardware failure. This causes vSAN to set the health status for all disks in the group to

Permanent disk loss, indicating disk failure.

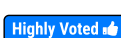
Assuming all other disks have not suffered from a hardware failure, why would vSAN mark all disks in the group as failed?

- A. The vSAN disk management service has failed.
- B. The affected vSphere host is offline.
- C. The key management server is offline.
- D. Deduplication and compression are enabled on the vSAN cluster.

Suggested Answer: A

Community vote distribution

D (100%)

  **adelbelkis2**  5 years, 1 month ago

D is the right answer

<https://kb.vmware.com/s/article/2149067>

upvoted 8 times

  **DenZn**  4 years, 11 months ago

Agree D is correct - Information from vSAN Deploy and Manage eBook: If an unhealthy disk belongs to a deduplication-enabled diskgroup, the whole diskgroup is marked as unhealthy.

upvoted 6 times

  **YoussefELL**  1 year, 10 months ago

Selected Answer: D

When deduplication and compression is enabled, if a capacity disk fails, the entire disk group becomes unavailable.

upvoted 1 times

  **Lazylinux** 3 years, 5 months ago

Selected Answer: D

D is correct



upvoted 1 times

  **ZakirK** 3 years, 5 months ago

Selected Answer: D

VMware Documentation

upvoted 1 times

  **diegof1** 4 years, 4 months ago

D is the correct answer.

Unhealthy disks or diskgroups are marked as such, and at this point, the disks or diskgroups are no longer used for new data placement. If an unhealthy disk belongs to a deduplication-enabled diskgroup, the whole diskgroup is marked as unhealthy.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - Unhealthy Devices and Data Evacuation section

upvoted 3 times

  **jasovj** 4 years, 10 months ago

D is correct, it's one of the caveats of dedup & compression, only on disk can fail wht whole diskgroup

upvoted 3 times

  **apiedra** 5 years ago

D is right, when you have dedup enabled and any disk failed, the entire diskgroup will go down

upvoted 5 times

  **Mzoear** 5 years ago

D is right

upvoted 2 times

In stretched clusters, what is a benefit of implementing a virtual witness rather than a physical witness?

- A. Reduced vSphere licensing
- B. Shared metadata between separate clusters
- C. Increased vSAN datastore capacity
- D. Increased compute for running VMs

Suggested Answer: A

Reference:

<https://cormachogan.com/2015/09/11/a-closer-look-at-the-vsan-witness-appliance/>

🗨️ 👤 **Lazylinux** 3 years, 5 months ago

A is answer

upvoted 1 times

🗨️ 👤 **diegof1** 4 years, 4 months ago

A is right

The witness host is located at the witness site and stores witness components for VM objects:

- Each stretched vSAN cluster must have its own witness host.
- The witness host is not a cluster member.

The host can be a physical host or a virtual appliance downloaded from the VMware website:

- Physical witness hosts:
 - Require standard vSphere licensing
 - Cannot start VMs that were previously running on the host
- Virtual witness hosts:
 - Are virtualized ESXi hosts running in a virtual appliance
 - Cannot run any VMs
 - Require less capacity and bandwidth than regular ESXi hosts
 - Are packaged by VMware with their own embedded license

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - About Witness Hosts section

upvoted 2 times

🗨️ 👤 **lotso** 4 years, 8 months ago

A. no doubt

upvoted 4 times

🗨️ 👤 **komodo08** 4 years, 8 months ago

A. Reduced vSphere Licensing

<https://storagehub.vmware.com/t/vsan-stretched-cluster-guide/witness-host-sizing-13/>

vSAN Witness Appliance (Virtual Machine)

Deploying the vSAN Witness Appliance that is provided by VMware is the recommended deployment choice for a vSAN Witness Host. When choosing this deployment option, there are some requirements to consider.

Licensing

A license is hard coded in the vSAN Witness Appliance and is provided for free from VMware.

Physical Host as a vSAN Witness Host

If using a physical host as the vSAN Witness Host there are some requirements to consider.

Licensing

If using a physical host as a vSAN Witness Host, it must be licensed with a valid vSphere license.

upvoted 4 times

🗨️ 👤 **jasonv** 4 years, 10 months ago

correct, no esxi license needed for new witness
upvoted 3 times

A storage administrator discovers vSAN is rebalancing components across a cluster randomly and faces degraded performance on the applications.

What should be done to avoid these issues?

- A. Size at least two disk groups on each node
- B. Maintain a minimum of 10 percent unused capacity
- C. Keep total storage consumption <70 percent
- D. Ensure there is sufficient queue depth on the I/O Controller

Suggested Answer: A


Reference:

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vsan-troubleshooting-reference-manual.pdf>

(24)

Community vote distribution

C (100%)

 **mariogg85** Highly Voted 5 years, 1 month ago

I believe the correct answer would be C

Note:

To provide enough space for maintenance and re protection, and to minimize automatic rebalancing events in the Virtual SAN cluster, consider keeping 30-percent capacity available at all times.


<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID-2EC7054E-FBCC-4314-A457-3DCAEDBCBD32.html>
upvoted 6 times

 **Lazylinux** Most Recent 2 years, 9 months ago

Selected Answer: C

C is the answer

upvoted 1 times

 **hsezer** 3 years, 1 month ago


I think A is correct. Because you already had a capacity issue and there is no chance to keep usage at %70. It is a design consideration, not a daily solution. Also to keep usage at %70 you should add additional capacity.

upvoted 1 times

 **Lazylinux** 3 years, 5 months ago

C is correct, disk groups is NOT mandatory how many...recently i had technical call wih VMware premium support with vSAN and during the discussion it was clearly pointed out that 25-30% minimum is required free space to ensure healthy vSAN and no bouncing activity of objects and components

upvoted 1 times


 **diegof1** 4 years, 4 months ago

C is the correc answer.

VMware recommends keeping a minimum of 25-30% of the vSAN datastore free. This recommendation is in-line with other HCI deployments. vSAN slack space ensures that vSAN has adequate space to perform functions such as:


- Maintenance mode data evacuation
- Resynchronization after a service interruption
- Application of a storage policy change that changes capacity needs
- Standard background activities such as the balancing of data across drives

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - vSAN Slack Space section
upvoted 1 times

 **rajVCIX** 4 years, 6 months ago



What about A for the answer? If you add another disk group it will also help with performance of applications wouldn't it?

upvoted 1 times

  **Virtual_ITTech** 3 years, 5 months ago


Performance degradation is only a symptom, the Root cause is that something causing the constant rebalancing due to vSAN Datastore is near full capacity

upvoted 1 times

  **Houston** 4 years, 7 months ago

C is the correct one

upvoted 1 times

  **jsonv** 4 years, 10 months ago

agree C, slack space will be always 30%+ if not you can have problems

upvoted 2 times

  **Mimmo2020** 4 years, 11 months ago

the correct answer is C

upvoted 1 times

  **Mzoear** 5 years ago

C is the correct answer

upvoted 1 times

What is the purpose of slack space in vSAN?

- A. Increases overall vSAN capacity
- B. Allows for maintenance and failure operations
- C. Stores vSphere HA heartbeats
- D. Provides vSAN analytic data

Suggested Answer: B

Reference:

<https://storagehub.vmware.com/t/vmware-vsan/vmware-r-vsan-tm-design-and-sizing-guide-2/how-much-slack-space-should-i-leave-1/>

Community vote distribution

B (100%)

🗲️ 👤 **jasonv** Highly Voted 👍 4 years, 10 months ago
B is correct slack space is used by background operations
upvoted 7 times

🗲️ 👤 **Lazylinux** Most Recent 🕒 3 years, 4 months ago
Selected Answer: B
is correct
upvoted 1 times

🗲️ 👤 **Lazylinux** 3 years, 5 months ago
B is right
upvoted 1 times

🗲️ 👤 **diegof1** 4 years, 4 months ago
B is right.

VMware recommends keeping a minimum of 25-30% of the vSAN datastore free. This recommendation is in-line with other HCI deployments.

vSAN slack space ensures that vSAN has adequate space to perform functions such as:

- Maintenance mode data evacuation
- Resynchronization after a service interruption
- Application of a storage policy change that changes capacity needs
- Standard background activities such as the balancing of data across drives

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - vSAN Slack Space

upvoted 1 times

The objects on a 4-node vSAN cluster are assigned a RAID-5 policy. A network outage occurs causing host one to lose connectivity with the rest of the cluster.

Seventy-five minutes elapse.

What is the health state of the objects?

- A. Reduced availability with no rebuild
- B. Reduced availability with no rebuild "" delay timer
- C. Non-availability related incompliance [non-compliance]
- D. Reduced availability

Suggested Answer: D

Community vote distribution

A (100%)

🗳️ 👤 **Lazylinux** 3 years, 4 months ago

Selected Answer: A

Is correct

upvoted 2 times

🗳️ 👤 **diegof1** 4 years, 4 months ago

A is the correct answer.

For RAID 5 needs $2n+2$, assuming FTT=1 because just exist 4 hosts. Host required are exactly 4 hosts to protect that failure.

Reduced availability with no rebuild: The object has suffered a failure, but VSAN was able to tolerate it. For example: I/O is flowing and the object is accessible. However, VSAN is not working on re-protecting the object. This is not due to the delay timer (reduced availability - no rebuild - delay timer) but due to other reasons. This could be because there are not enough resources in the cluster, or this could be because there was not enough resources in the past, or there was a failure to re-protect in the past and VSAN has yet to retry. Refer to the limits health check for a first assessment if any resources may be exhausted. You have to resolve the failure or add resources as quickly as possible in order to get back to being fully protected against a subsequent failure.

<https://kb.vmware.com/s/article/2108319>

upvoted 2 times

🗳️ 👤 **jasonv** 4 years, 10 months ago

A is correct

upvoted 2 times

🗳️ 👤 **LCOJ** 4 years, 11 months ago

I agree with A as per vSAN desing, RAID 5 uses 4 nodes to be compliance.

upvoted 2 times

🗳️ 👤 **Mzoear** 5 years ago

B is the correct answer

upvoted 1 times

🗳️ 👤 **Mzoear** 5 years ago

sorry i think A is correct

upvoted 2 times

🗳️ 👤 **mariogg85** 5 years ago

IMHO should be A

Arguments: Minimum number of hosts required for Raid 5 are $3 + 1$. Then with 4 nodes vSAN Cluster if one host fails vSAN couldn't rebuild because the component has lost the parity.

Reference: <https://kb.vmware.com/s/article/2108319>

Reduced availability with no rebuild: The object has suffered a failure, but VSAN was able to tolerate it. For example: I/O is flowing and the object is

accessible. However, VSAN is not working on re-protecting the object. This is not due to the delay timer (reduced availability - no rebuild - delay timer) but due to other reasons. This could be because there are not enough resources in the cluster, or this could be because there was not enough resources in the past, or there was a failure to re-protect in the past and VSAN has yet to retry. Refer to the limits health check for a first assessment if any resources may be exhausted. You have to resolve the failure or add resources as quickly as possible in order to get back to being fully protected against a subsequent failure.

Arguments:

upvoted 2 times

  **anujspoddar** 5 years, 1 month ago

Should be A,

Minimum number of hosts required for Raid 5: $2n+1$

Reduced availability - active rebuild: The object has suffered a failure, but it was configured to be able to tolerate the failure. I/O continues to flow and the object is accessible. vSAN is actively working on re-protecting the object by rebuilding new components to bring the object back to compliance.

Reduced availability with no rebuild: The object has suffered a failure, but VSAN was able to tolerate it. For example: I/O is flowing and the object is accessible. However, VSAN is not working on re-protecting the object. This is not due to the delay timer (reduced availability - no rebuild - delay timer) but due to other reasons. This could be because there are not enough resources in the cluster, or this could be because there was not enough resources in the past, or there was a failure to re-protect in the past and VSAN has yet to retry.

upvoted 2 times

  **adelbelkis2** 5 years, 1 month ago

A

<https://kb.vmware.com/s/article/2108319>

upvoted 3 times

A vSAN 5-node cluster has two sets of VMs, each associated with a separate storage policy.

-Group A (Storage Policy: FTT=0)

-Group B (Storage Policy: FTT=2)


What are two possible results when two hosts in the vSAN cluster fail permanently? (Choose two.)

- A. VMs in Group B will become inaccessible if a third host permanently fails.
- B. vSAN will rebuild data associated with Group A.
- C. VMs in Group A might experience data inaccessibility.
- D. VMs in Group B might experience data loss.
- E. A host can be put into maintenance mode without impacting VMs in either Group A or B.

Suggested Answer: CD

Community vote distribution

AC (100%)

 **mariogg85** Highly Voted 5 years, 1 month ago

I think that the correct answers may be A and C.

Answer E: It's not possible because VMs Group A with FTT=0 not support any host failure. Then VMs Group A couldn't support 3 Host Failure or Maintenance.

Answer B : VMs Group A not support 2 Host failure. Then vSAN cannot rebuild

Answer D: VMs Group B should support with FTT = 2 (n+1). Then with 2 Host on permanent fail the VMs Group B not experience data loss.
upvoted 9 times

 **jasonv** Highly Voted 4 years, 10 months ago

i'll go with A and C as mariogg85 explained
upvoted 5 times

 **hsezer** Most Recent 3 years, 1 month ago

Lets think about options. Not very clear scenario but my opinion is A & C are most suitable ones.

a- You can not be sure if two copies were on the failed hosts or rebuild is completed. So it is an probability but not certain case.

b- Group A is FTT=0, so it is stored as 1 copy. So if data has lost there is no chance to rebuild.

c- If the objects of a vm is failed hosts yes you might experience inaccessibility.

d- FTT=2 with 5 hosts mean raid-1 with 3 copies. If 2 copies are on the failed hosts still there is one copy, so it is not possible to lose data.

e- After 2 hosts failure, there are 3 hosts survive. Think 3 copies of Group-B are there and put the one host in maintenance impact compliance of FTT=2.

upvoted 1 times

 **Jonesythegreat** 3 years, 4 months ago

My vote is C and D. For sure there would be some objects from Group A with data loss because they were on the 2 failed hosts. Group B supports two failures and with two hosts dropping would have the opportunity for some objects to have data loss as well. I think C and D would be the most correct.

upvoted 1 times

 **Lazylinux** 3 years, 5 months ago

Selected Answer: AC


Are correct

upvoted 1 times

 **Lazylinux** 3 years, 5 months ago

I vote A & C

upvoted 1 times

 **diegof1** 4 years, 4 months ago

A & C is the correct answer.

Group A doesn't support any failure, while Group B just support two. and if a third host fail data would be inaccessible.

upvoted 2 times

  **DSA** 4 years, 11 months ago

Tricky question with "sub-optimal" options for the Answer...

A => Some VMs could be lucky to have their last two copies (non-witness objects) on the surviving host. so the Data is there and not lost. But why should all of Group B become inaccessible? Because of the missing ability to build a quorum?

B => yeah, that can be ruled out pretty surely.

C => not sure about that, because of "might experience". The VMs in Group A will definitely have a data loss (not just inaccessible) - if not already suffered before. And lost data is at the same time inaccessible (because it is lost). But the wording just does not "feel right".

D => similar to A - the "unlucky" VMs have a data loss (just the witness-objects are left).

E => with just two host left? That can be ruled out too.

Personally I would go for A & D:

A => Quorum can not be established => all of Group B are inaccessible.

D => Some of the (inaccessible) VMs could also have suffered Data-Loss on top of that (only their witness-objects are left on the last 2 surviving hosts, all mirrors are gone)

upvoted 1 times

  **mw100** 4 years, 11 months ago

D cannot be, as it will tolerage 2 failures. A&C is correct answers.

upvoted 4 times

A 3-node vSAN cluster will take which two actions, when a drive is physically removed from a vSAN node? (Choose two.)

- A. Marks all components on that device as absent
- B. Marks all components on that device as degraded
- C. Waits for the configured delay timer before rebuild
- D. Marks all components on that device as stale
- E. Starts component rebuild immediately

Suggested Answer: AC

Reference:

<https://storagehub.vmware.com/t/vmware-vsan/vsan-6-7-proof-of-concept-guide/disk-failures-9/>

Community vote distribution

AC (100%)

🗲️ 👤 **Lazylinux** 3 years, 5 months ago

Selected Answer: AC

<https://cormachogan.com/2014/12/04/vsan-part-30-difference-between-absent-degraded-components/>

upvoted 1 times

🗲️ 👤 **Lazylinux** 3 years, 5 months ago

A & C is the answer

upvoted 1 times

🗲️ 👤 **diegof1** 4 years, 4 months ago

A & C is correct.

If a capacity device fails without an error, the cluster marks all components on that device as absent and waits the configured delay before beginning to rebuild the affected components.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - Failure Scenario: Rebuilding Components section

upvoted 1 times

🗲️ 👤 **jasonv** 4 years, 10 months ago

yes, as we pull off the disk, vsan don't mark object as inaccessible, it waits 60 minutes timer before begin the rebuild, if possible.

upvoted 3 times

🗲️ 👤 **mariogg85** 4 years, 12 months ago

I'm agree with the answer

upvoted 3 times

A vSAN cluster has this configuration:

- 4 hosts with 1 disk group per host
- Each disk group contains 1 cache device and 7 capacity devices

What are two ways to increase the vSAN datastore storage capacity? (Choose two.)

- A. Add a host to the cluster
- B. Add a cache device to each disk group
- C. Add a capacity drive to each disk group
- D. Replace a cache device with a larger cache device
- E. Add a disk group to each host

Suggested Answer: AE

Community vote distribution

AE (100%)

🗲️ 👤 **jasonv** Highly Voted 👍 4 years, 10 months ago

AE is correct

upvoted 7 times

🗲️ 👤 **mariogg85** Highly Voted 👍 4 years, 12 months ago

Agree. Right answers

upvoted 5 times

🗲️ 👤 **Lazylinux** Most Recent 🕒 3 years, 4 months ago

Selected Answer: AE

Is correct

upvoted 1 times

🗲️ 👤 **Lazylinux** 3 years, 5 months ago

A & E for sure

upvoted 1 times

🗲️ 👤 **AllenHuang** 4 years, 1 month ago

Agree, A&E

upvoted 1 times

🗲️ 👤 **diegof1** 4 years, 4 months ago

A & E are right.

upvoted 1 times

🗲️ 👤 **guruto** 4 years, 5 months ago

Correct A,E

no more capacity disk is able to add on the DGs.

<https://configmax.vmware.com/guest?vmwareproduct=vSphere&release=vSphere%206.7&categories=7-0>

upvoted 3 times

An organization is no longer compliant with security requirements because a vSphere administrator disabled encryption on a vSAN cluster. Which vCenter Server role ensures only authorized vSphere administrators have access to encryption functionality?

- A. Administrator
- B. Read-only
- C. No cryptography administrator
- D. No access

Suggested Answer: C

Reference:

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.security.doc/GUID-6C181D08-6650-4AD1-92D1-AAFDA3A3E38C.html>

🗳️ 👤 **Lazylinux** 3 years, 5 months ago

Agree with Diegof1 explanation and Answer is A..the Key to the question is "Have Access to Encryption" and hence cannot be C
upvoted 2 times

🗳️ 👤 **HRC** 3 years, 7 months ago

Agree with DenZn and DSA. Answer should be A. Reason for not opting C is that 'no Cryptography Administrator' is not permitted to Encryption operations.
upvoted 1 times

🗳️ 👤 **vJoeyB** 4 years ago

C is right - the question asks which role ensures only authorized admins have access to encrypt. If you assign everyone the admin role, this does not ensure that only authorized admins have access to encrypt.

If you leverage the no cryptography admin role, then you are ensuring that only authorized admins have access to encrypt.

upvoted 2 times

🗳️ 👤 **Stekman** 4 years, 4 months ago

I think the answer is C. The existing of this role allow you to have administrators with no encryption rights.
If this role did not exists every administrator has the capability.
upvoted 1 times

🗳️ 👤 **diegof1** 4 years, 4 months ago

C is wrong: Users with the No cryptography administrator role for an object have the same privileges as users with the Administrator role, except for Cryptographic operations privileges. This role allows administrators to designate other administrators that cannot encrypt or decrypt virtual machines or access encrypted data, but that can perform all other administrative tasks.

On the other hand, Only administrators with encryption privileges can perform encryption and decryption tasks. I think the correct answer is A.

upvoted 2 times

🗳️ 👤 **alsmk2** 4 years, 9 months ago

This is a really badly worded question - could be either A or C depending on how you read it.

I'd side with C as that would prevent access to encryption functionality, and it's fairly new, so would make sense that they would ask about it.

upvoted 1 times

🗳️ 👤 **DSA** 4 years, 6 months ago

Yeah, it is a little bit confusing. BUT: If the wording is 1:1 exact, then it is A.

=> "Which role >ensures< ... >have access to encryption<"

upvoted 4 times

🗳️ 👤 **DenZn** 4 years, 10 months ago



The correct anwer should be A

<https://storagehub.vmware.com/t/vsan-6-7-update-1-technical-overview/role-based-access-control-4/>

vSphere 6.5 introduced the No Cryptography Administrator role along with the introduction of VM Encryption. This role is very similar to the normal

administrator with many of the same privileges. Operations such as power on or off a virtual machine, boot, shutdown, vMotion, as well as normal vSAN management may be performed. However, this role is not allowed to perform any cryptographic operations.

upvoted 4 times

  **jasonv** 4 years, 10 months ago

C is correct, its a new role.

upvoted 1 times

vCenter Server is offline.

What are two ways to check the health of a vSAN cluster? (Choose two.)

- A. ESXi system logs on vSAN datastore
- B. vSphere Host Client
- C. esxcli
- D. HCI Bench
- E. esxtop

Suggested Answer: BC

Reference:

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.security.doc/GUID-6C181D08-6650-4AD1-92D1-AAFDA3A3E38C.html>

Community vote distribution

BC (100%)

 **Lazylinux** 3 years, 4 months ago

Selected Answer: BC


correct

upvoted 1 times

 **Lazylinux** 3 years, 5 months ago

B and C for sure

upvoted 1 times

 **diegof1** 4 years, 4 months ago

B&C are correct.

With vSAN 6.6 and newer versions, we can point a web browser directly at the ESXi management IP or FQDN of any host in the vSAN cluster and log in to the vSphere Host Client to monitor vSAN. As seen in the screenshot below, many of the vSAN Health items are visible in this UI.

<https://blogs.vmware.com/virtualblocks/2018/04/05/vsan-when-vcenter-server-is-offline/>

Use Esxcli commands to obtain information about vSAN and to troubleshoot your vSAN environment.

The following commands are available:

Command ----- Description

esxcli vsan network list ----> Verify which VMkernel adapters are used for vSAN communication.

esxcli vsan storage list ----> List storage disks claimed by vSAN.


esxcli vsan cluster ----> get Get vSAN cluster information.

esxcli vsan health ----> Get vSAN cluster health status.

esxcli vsan debug ----> Get vSAN cluster debug information.

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vsan-monitoring.doc/GUID-7799D2D7-2513-4372-92EA-4A1FB510E012.html>

upvoted 1 times

 **Houston** 4 years, 7 months ago

I agree as well

upvoted 2 times

 **anujspoddar** 5 years, 1 month ago

Better reference:

<https://blogs.vmware.com/virtualblocks/2018/04/05/vsan-when-vcenter-server-is-offline/>

upvoted 4 times

  **mariogg85** 4 years, 12 months ago

I'm agree with you. It's the better link reference. The answers are right.
upvoted 2 times

What are two purposes of a vSAN storage policy? (Choose two.)

- A. Determine how storage objects are provisioned
- B. Determine vSAN encryption level
- C. Guarantee the required level of service
- D. Enable deduplication and compression
- E. Enable TRIM/UNMAP

Suggested Answer: AC

Reference:

<http://www.vmwarearena.com/understanding-vmware-virtual-san-storage-policies-detailed-view/>

Community vote distribution

AC (100%)

LC0J **Highly Voted** 4 years, 11 months ago

In my opinion it's right
upvoted 7 times

jasonv **Highly Voted** 4 years, 10 months ago

yes, ruling out the other answers it's correct
upvoted 7 times

RRK06 **Most Recent** 3 years, 1 month ago

Selected Answer: AC
A&C are correct indeed
upvoted 1 times

Lazylinux 3 years, 4 months ago

Selected Answer: AC
is correct
upvoted 1 times

Lazylinux 3 years, 5 months ago

A & C for sure
upvoted 1 times

AllenHuang 4 years, 1 month ago

Agree, A&C
upvoted 2 times

diegof1 4 years, 4 months ago

A & C are correct.
upvoted 2 times

An administrator wishes to prevent notifications of vSAN Health status during a scheduled maintenance window.
What can be done to achieve this?

- A. Disable SNMP service
- B. Disable the performance service
- C. Run performance diagnostics prior to scheduled maintenance
- D. Silence the alert from the HTML client

Suggested Answer: A

Community vote distribution

D (100%)

🗳️ 👤 **MSK12** Highly Voted 4 years, 11 months ago

Yeah, D is correct
upvoted 7 times

🗳️ 👤 **Mzoear** Highly Voted 5 years ago

i think D is the right answer
upvoted 6 times

🗳️ 👤 **Lazylinux** Most Recent 3 years, 4 months ago

Selected Answer: D
IS D is good
upvoted 1 times

🗳️ 👤 **Lazylinux** 3 years, 5 months ago

Agree D is answer
upvoted 1 times

🗳️ 👤 **vJoeyB** 4 years ago

D is correct <https://kb.vmware.com/s/article/2151813>
upvoted 3 times

🗳️ 👤 **sheriefmb** 4 years, 5 months ago

Answer is A as there is no way to silence the alert from HTML client. we can only do it via RVC and API
<https://kb.vmware.com/s/article/2151813>
upvoted 3 times

🗳️ 👤 **lollo1234** 4 years, 4 months ago

That's for vSAN 6.6. From vSAN 6.7 there is an option on the GUI.
<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vsan-monitoring.doc/GUID-9504EECF-5946-49FB-86C6-8A4F977F5FC3.html>
upvoted 5 times

An administrator is planning to deploy a new vSAN cluster with these requirements:

- Physical adapters will share capacity among several traffic types
- Guaranteed bandwidth for vSAN during bandwidth contention
- Enhanced security and performance

What must be considered when configuring the new vSAN cluster?

- A. Isolate traffic in a VLAN and utilize Network I/O Control
- B. Utilize IOPS Limit rules in storage policies
- C. Utilize jumbo frames with Network I/O Control
- D. Implement vSphere I/O filters

Suggested Answer: A

Reference:

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vsan-planning.doc/GUID-031F9637-EE29-4684-8644-7A93B9FD8D7B.html>

Community vote distribution

A (100%)

 **Lazylinux** 3 years, 5 months ago

Selected Answer: A


A is correct

upvoted 1 times

 **Lazylinux** 3 years, 5 months ago

A for sure

upvoted 1 times

 **diegof1** 4 years, 4 months ago

A is the correct answer.

Consider the following when preparing a vSAN network:

- For best networking performance, all hosts must be connected in the same subnet. In vSAN 6.0 and later, you can also connect hosts in the same layer 3 network if necessary.
- Dedicate network bandwidth on a physical adapter and allocate at least 1 Gbps of bandwidth for vSAN. You might use one of the following configuration options:
 - Dedicate 1-GbE physical adapters for a hybrid host configuration.
 - Use dedicated or shared 10-GbE physical adapters for all-flash configurations.
 - Use dedicated or shared 10-GbE physical adapters for hybrid configurations if possible.
 - Direct vSAN traffic on a 10-GbE physical adapter that handles other system traffic and use Network I/O Control on a distributed switch to reserve bandwidth for vSAN.
- Configure a port group on a virtual switch for vSAN:
 - Assign the physical adapter for vSAN to the port group as an active uplink.
 - You can also assign vSAN traffic to a VLAN by enabling tagging in the virtual switch.

Taken from vSAN 6.7 Deploy And Manage Lecture Manual

upvoted 2 times

 **jasonv** 4 years, 10 months ago

yes, ruling out, we must create a dedicated workgroup for vsan and use NIOC in case there is contention on the network. Using jumbo frames can be another option but sharing vsan traffic in a workgroup together with other types it's a no-no

upvoted 3 times

Which two conditions should be verified before removing the ESXi host from a vSAN cluster? (Choose two.)

- A. Data evacuation is complete.
- B. Encryption is disabled.
- C. ESXi host is in maintenance mode.
- D. Resyncs are running.
- E. All objects are currently healthy.

Suggested Answer: CE

Reference:

<https://vsan-essentials.gitbooks.io/vsan-6-2/content/chapter7.html>

🗨️ **wtkc** Highly Voted 4 years, 9 months ago

We can choose "no evacuation" option. so, A is not necessary.

Actually, C and E are correct answers.

upvoted 7 times

🗨️ **DSA** 4 years, 6 months ago

Would agree to C & E:

When the host is in Maintenance, the evacuation operation is completed (->A). E is obvious ;)

upvoted 2 times

🗨️ **dfdd1** 3 years, 10 months ago

you can put the host in Maintenance mode without data evacuation. So it doesn't mean evacuation is done or not.

upvoted 3 times

🗨️ **Virtual_ITTech** 3 years, 5 months ago

you have to have data evacuation, the question is talking about completely removing a host from cluster, and not just a temporary host shutdown. so A is must have.

upvoted 1 times

🗨️ **Ady_14** Highly Voted 4 years, 11 months ago

I agree with A and C.

<https://kb.vmware.com/s/article/2148975>

upvoted 7 times

🗨️ **hsezer** Most Recent 3 years, 1 month ago

I think A is not necessary, you can put the host in maintenance mode also selecting ensure accessibility. After object repair timer delay time objects will be synchronized. So the key point is healthy status. A similar operation of removing disk groups says you can use also ensure accessibility option.

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/GUID-16EBFE47-28BE-48DA-8B62-C99B2A7DC5C0.html>

upvoted 1 times

🗨️ **hsezer** 3 years, 1 month ago

C and E correct.

upvoted 1 times

🗨️ **Lundu1995** 3 years, 4 months ago

In my opinion it's A & E because the best way is to evacuate the host (full data migration) and check if every object is happy. then remove the node. for sure the node must be in maintenance, but in my opinion this is obvious. I guess they want to hear which option you want to choose and this is full data migration and check if everything is happy

upvoted 1 times

🗨️ **Virtual_ITTech** 3 years, 5 months ago

I would say A, E. A must have because data evacuation is required when permanently removing a host from cluster. E is more likely than C because we want to make sure objects are healthy, which is the same thing as waiting for resync to complete. C is likely but it is obvious because putting a host in MM is inherent to the process of decommissioning a host and specific to vSAN situation

upvoted 2 times

🗨️ 👤 **Lazylinux** 3 years, 5 months ago

A & C for sure..remember here you are removing completely Esxi Host from cluster and hence if NO complete evacuation done then you will end up with orphaned or inaccessible VMS

upvoted 2 times

🗨️ 👤 **Lazylinux** 3 years, 5 months ago

I correct myself since i been reading more and i would go for A and E => because before evacuating data one MUST make sure all objects are Healthy then FULL data evacuation, because when you try to remove host from vSAN it will ask you about Modes and hence you already getting into maintenance mode and hence C no really applicable

upvoted 1 times

🗨️ 👤 **HRC** 3 years, 6 months ago

I would go for A and E on as a 'verification' before removing a host from vSAN cluster. Removing a host will remove the disk group(s) so is the data on it as well.

upvoted 1 times

🗨️ 👤 **Char250** 3 years, 8 months ago

Correct answers are A & C <https://kb.vmware.com/s/article/2148975>

2. Put the host into maintenance mode with full data migration selected. This will evacuate all disk groups.

upvoted 2 times

🗨️ 👤 **atti9** 3 years, 10 months ago

You migrate the data out of the node by putting it into maintenance mode. E is a good practice, but objects may be unavailable for other unrelated reasons that does not hinder the node from being removed. So A and C are correct.

upvoted 1 times

🗨️ 👤 **RoninVI** 3 years, 11 months ago

Right-click on host > Select Maintenance Mode > Select data handling option > click Ok > Data Evac Completes > Resynch objects completes > Host is now in Maintenance Mode.

upvoted 2 times

🗨️ 👤 **diegof1** 4 years, 4 months ago

A & C are correct.

You need to put the host in maintenance mode with the recommended option "Full Data migration" then wait for data evacuation to be complete.

Put a host in maintenance mode is not a guarantee that data is evacuated because you could choose "No data migration". This is not a good practice, because you should ensure the data is available in the rest hosts before remove it from the vSAN Cluster.

upvoted 4 times

🗨️ 👤 **Eashwar** 4 years, 6 months ago

A and C should be the correct answer .

The option E is the general verification which need to be done for any maintenance activity.

upvoted 1 times

🗨️ 👤 **rajVCIX** 4 years, 6 months ago

I would go with C because you do not need to evacuate the data from the host before moving it from the cluster. You can select the option to ensure availability. The question didn't give enough information to suggest that A is the only feasible answer which is data evacuation.

E to ensure that the objects are healthy before a host cluster removal.

upvoted 4 times

🗨️ 👤 **Houston** 4 years, 7 months ago

I would go for A and C

upvoted 2 times

🗨️ 👤 **alsmk2** 4 years, 9 months ago

A & E -

A: Data evacuation would suggest the host is already in maintenance mode, making C redundant.

E: Would make sense that all objects should be healthy, with valid storage policies before removing the host.

upvoted 3 times

🗨️ 👤 **MIP** 4 years, 9 months ago

personally I would look for A & E before I remove a host.

A "data evacuation is complete" means that there are no more resyncs running & most probably that the host is in maintenance mode or that most

probably all disk groups of the hosts are removed (= data evacuated). that is important before you remove a host!


B is non sense.

C is technically possible without evacuating data.

D resyncs are not running anymore when all objects are healthy and the data evacuation is complete.

E all objects should be healthy before you remove a host!

upvoted 2 times

  **LC0J** 4 years, 11 months ago

A agree that correct answers should be A and C as per explanation below:

<http://www.vexpertconsultancy.com/2019/06/step-by-step-remove-a-node-permanently-from-a-vsan-cluster/>

upvoted 5 times

Which VM file type resides in the VM home namespace object on a vSAN datastore?

- A. .vswp
- B. .vmsn
- C. .vmx
- D. .vmem

Suggested Answer: C

Reference:

<https://vsan-essentials.gitbooks.io/vsan-6-2/content/chapter6.html>

Community vote distribution

C (100%)

🗳️ 👤 **Lazylinux** 3 years, 4 months ago

Selected Answer: C

is correct

upvoted 1 times

🗳️ 👤 **Lazylinux** 3 years, 5 months ago

C for sure

upvoted 1 times

🗳️ 👤 **diegof1** 4 years, 4 months ago

C is the correct answer

When you provision a virtual machine on a vSAN datastore, a set of objects is created. These objects are of the following types:

- VM home namespace: Stores the virtual machine metadata (configuration files).
- VMDK: Virtual machine disk.
- VM swap file: Virtual machine swap file which is created when the virtual machine is powered on.
- Snapshot deltas: Created when snapshots of the virtual machine are taken.
- VM memory: Virtual machine memory state when a virtual machine is suspended or when the snapshot of a virtual machine is taken and its memory state is preserved.

Virtual machine files that reside in the VM home namespace include:

- The .log and .hlog files
- The .nvram file
- The .vmsd file
- The .vmx file

Taken from vSAN 6.7 Deploy and Manage Lecture Manual

upvoted 2 times

🗳️ 👤 **alsmk2** 4 years, 9 months ago

Correct Answer

<https://cormachogan.com/2014/03/12/vsan-part-18-vm-home-namespace-and-vm-storage-policies/>

upvoted 3 times

In a 10-node hybrid vSAN cluster, each node has seven 4TB magnetic disks and one 2TB SSD.

What is the raw storage capacity of the vSAN datastore?

- A. 140TB
- B. 150TB
- C. 280 TB
- D. 300 TB

Suggested Answer: C

🗲️ 👤 **LCOJ** Highly Voted 4 years, 11 months ago

The answer C is correct if you think only about Capacity where we have 10 nodes with 7 disks with 4 TB, this way $10 * 28 \text{ TB}$ per node = 280 TB raw to be used by vSAN Cluster

upvoted 8 times

🗲️ 👤 **MVSM27** Highly Voted 4 years, 7 months ago

C is the correct because if you consider a Hybrid cluster using 7 magnetic disks + 1 SSD disk for each node, the SSD is only for cache not for capacity. 7 disks (4TB each) x 10 nodes = 280 TB

upvoted 6 times

🗲️ 👤 **Lazylinux** Most Recent 3 years, 5 months ago

C for sure

upvoted 1 times

🗲️ 👤 **diegof1** 4 years, 4 months ago

C is right.

Total Capacity of host = $7 \times 4 \text{ TB} = 28 \text{ TB}$

Total Capacity of Cluster = $10 \times 28 \text{ TB} = 280 \text{ TB}$

upvoted 2 times

🗲️ 👤 **jasonv** 4 years, 10 months ago

yes 280, C is correct

upvoted 5 times

🗲️ 👤 **LCOJ** 4 years, 11 months ago

The answer A is correct if you think only about Capacity where we have 10 nodes with 7 disks with 4 TB, this way $10 * 28 \text{ TB}$ per node = 280 TB raw to be used by vSAN Cluster

upvoted 2 times

🗲️ 👤 **desolator** 4 years, 4 months ago

Your math is right, but I am sure you meant to say the answer is C.

upvoted 2 times

Refer to the exhibit.

	vSAN Node	vSAN Cluster
Nodes/Arrays Required	1	4
Capacity Disks	15	60
Flash Caching Devices	3	12
Disk Groups	3	12
Capacity disks per Disk Group	5	5
Total Raw Capacity (TB)	36	144
Total Usable Capacity (TB)	18	72
Total Flash Cache Capacity (TB)	1.2	3.6

Components	Qty	Details	Total
CPU	2	Intel "Xeon" Gold 6126 12C 125W 2.6GHz	2
MEM	4	ThinkSystem 32GB TruDDR4 2666 MHz (1Rx1.2V) RDIMM	128
Controller	1	ThinkSystem 430-8I SAS/SATA 12 Gb HBA	1
NIC	1	Intel Dual-Port 100Gbp/s	2
Flash Cache	3	ThinkSystem 2.5" HUSMM32 400GB Performance SAS 12GB Hot Swap SSD	1.2
Persistent Disk	15	ThinkSystem 2.5" 2.4TB 10K SAS 12GB Hot Swap 512e HDD	36

A solution architect is setting up a highly resilient vSAN hybrid cluster.

Why is the cache incorrectly sized in this configuration?

- A. Cache disk needs to meet vSAN Compatibility Guide
- B. Cache-to-capacity ratio needs to be sized lower
- C. Cache-to-capacity ratio needs to be sized higher
- D. Number of cache disks in disk group is insufficient.

Suggested Answer: D

Community vote distribution

C (100%)

 **DenZn**  4 years, 11 months ago

I think Answer C is correct.

As a guideline, VMware recommends having at least a 10% flash cache to consumed capacity ratio in Hybrid vSAN configurations.

If you have 36TB capacity, you will need 3,6TB cache - But there you only have 1,2 TB cache.

upvoted 11 times

🗨️ **ALLOY** 3 years, 10 months ago

The flash caching device must provide at least 10 percent of the anticipated storage that virtual machines are expected to consume, not including replicas such as mirrors.

upvoted 1 times

🗨️ **DenZn** 4 years, 11 months ago

<https://storagehub.vmware.com/t/vmware-r-vsan-tm-design-and-sizing-guide-2/cache-sizing-overview-1/>

upvoted 2 times

🗨️ **ALLOY** 3 years, 10 months ago

so it must be 1.8 rather than 1.2 , correct me if im wrong

upvoted 2 times

🗨️ **ALLOY** 3 years, 10 months ago

C is correct , but i think we should look at the used capacity not the raw capacity cuz it might be used for either storage and data protection

upvoted 2 times

🗨️ **LCOJ** 4 years, 11 months ago

In this case you have Only 1.2 TB cache because there are no required cache disks to reach 3 cache devices per node * 4 nodes to be compliance with the 10% cache required as per vSAN Design Guide

upvoted 1 times

🗨️ **Lundu1995** Most Recent 3 years, 4 months ago

The easiest way to calculate it is Diskgroup Raw capacity * 10% = Cache size. In this case the cache is too small

upvoted 1 times

🗨️ **Jonesythegreat** 3 years, 4 months ago

I vote C because it looks like they forgot to multiply capacity disks per host to get the cluster number. Cluster capacity is only 5 disks for capacity and is wrong.

upvoted 1 times

🗨️ **Lazylinux** 3 years, 5 months ago

Selected Answer: C

is correct

upvoted 2 times

🗨️ **Lazylinux** 3 years, 5 months ago

C is correct as per comments

upvoted 1 times

🗨️ **safodz** 4 years ago

total cache capacity by cluster should be 4.8 and not 3.6 as mentioned in the post

the correct answer is C

upvoted 1 times

🗨️ **GodMan114** 4 years, 4 months ago

B.

Cache to Capacity ratio needs to be lower. A "low ratio" means a larger-sized cache drive compared to capacity. I think it's a trick question, which is silly because this exam should be to ensure we know the content, not trick us with words.

Here is a reference to VMware calling it "high ratio" meaning you need more cache:

<https://core.vmware.com/resource/vmware-vsan-design-guide#sec4-sub6>

upvoted 2 times

🗨️ **diegof1** 4 years, 4 months ago

The correct answer is C.

Best practice: Ensure there is enough cache to meet the design requirements. The recommendation for cache is 10% of the anticipated consumed storage capacity before the NumberOfFailuresToTolerate is considered.

If the RAW Capacity for each host is 36TB then the cache should be 3.6 TB for each host.

<https://core.vmware.com/resource/vmware-vsan-design-guide#sec5-sub5>

upvoted 1 times

🗨️ 👤 **diegof1** 4 years, 4 months ago

The correct answer is C.

Best practice: Ensure there is enough cache to meet the design requirements. The recommendation for cache is 10% of the anticipated consumed storage capacity before the NumberOfFailuresToTolerate is considered.

<https://core.vmware.com/resource/vmware-vsan-design-guide#sec5-sub5>

upvoted 1 times

🗨️ 👤 **wtkc** 4 years, 9 months ago

I agree with what was written before.

but using larger cache disks means making cache-capacity ratio "lower", i think.

So, I choose B.

upvoted 3 times

🗨️ 👤 **jasonv** 4 years, 10 months ago

you have 12 DG per cluster and 12 flashing devices, so correct, but undersized cache disk

upvoted 1 times

🗨️ 👤 **jasonv** 4 years, 10 months ago

I thing C is correct, cache is undersized on that diskgroups, Vsan disk groups only need a cache disk correctly sized, but take into account that on all flash the limit of usable cache is 600GB

upvoted 1 times

🗨️ 👤 **LCOJ** 4 years, 11 months ago

I think the D is correct, once you have 4 nodes into vSAN Cluster you are going to need 4 cache devices, one per disk group.

upvoted 3 times

🗨️ 👤 **DenZn** 4 years, 11 months ago

You have 3 Diskgroups per Nodes and also have 3 Chace Devices per Node. For the Cluster you need 12 cache Devices.

upvoted 3 times

🗨️ 👤 **akfriend** 3 years, 6 months ago

Last row of first table states required cluster cache capacity is 3.6(equates to 9 disks) however with 12 cache drives of 400GB, it should be 4.8 so chache is incorrectly sized. Correct answer is D

upvoted 1 times

An all-flash vSAN cluster contains four nodes.

Which two storage policies can the cluster satisfy? (Choose two.)

- A. FTT=2 (RAID-1 Mirroring)
- B. FTT=2 (RAID-6 Erasure Coding)
- C. FTT=1 (RAID-1 Mirroring)
- D. FTT=1 (RAID-5 Erasure Coding)
- E. FTT=3 (RAID-1 Mirroring)

Suggested Answer: CD

Community vote distribution

CD (100%)

🗲️ 👤 **Lazylinux** 3 years, 4 months ago

Selected Answer: CD

is correct

upvoted 1 times

🗲️ 👤 **Lazylinux** 3 years, 5 months ago

C and D correct

upvoted 1 times

🗲️ 👤 **diegof1** 4 years, 4 months ago

C & D are right

The required number of hosts for RAID 1 is $2n+1$ while for RAID5/6 is $2n+2$, so with 4 host just can support FTT=1 in RAID 1 [$2(1)+1=3$ host] and FTT=1 in RAID 5 [$2(1)+2=4$ hosts]

upvoted 3 times

🗲️ 👤 **LC0J** 4 years, 11 months ago

I agree with answer C and D using RAID 1 and RAID 5 with FTT = 1

upvoted 3 times

🗲️ 👤 **mariogg85** 5 years, 1 month ago

Right answer

<https://docs.vmware.com/en/VMware-Cloud-on-AWS/services/com.vmware.vsphere.vmc-aws-manage-data-center-vms.doc/GUID-EDBB551B-51B0-421B-9C44-6ECB66ED660B.html>

upvoted 3 times

A company implements a vSAN environment with linked clone virtual desktops and homogeneous desktop operating systems. During the last three months, users have experienced intermittent high latency, degraded performance, and lockout from desktops. Which two methods should an administrator use to monitor and respond to the issues? (Choose two.)

- A. vRealize Log Insight
- B. vRealize Automation
- C. Live Optics
- D. HCI Bench
- E. vRealize Operations Manager

Suggested Answer: AC

Community vote distribution

AE (100%)

🗲️ 👤 **Mzoear** Highly Voted 5 years ago

May be A & E

upvoted 8 times

🗲️ 👤 **Lazylinux** Most Recent 3 years, 5 months ago

Selected Answer: AE

Is correct

upvoted 1 times

🗲️ 👤 **Lazylinux** 3 years, 5 months ago

A and E i believe are correct

upvoted 1 times

🗲️ 👤 **wtkc** 4 years, 9 months ago

Actually A and C and E are correct, i think.

but vROps is not mentioned in ebook as vsan interoperable software.

so, should be A and C.

upvoted 2 times

🗲️ 👤 **alsmk2** 4 years, 8 months ago

Live Optics is a a non-vmware product not used for troubleshooting. A & E are the only sensible choices here - vRealize has dashboard for vSAN, and integrates completely with it via the vSphere adapter.

upvoted 4 times

🗲️ 👤 **jasonv** 4 years, 10 months ago

A - E is correct

upvoted 3 times

🗲️ 👤 **LCOJ** 4 years, 11 months ago

The correct answer is A and E where you can view erros, logs and analytic informations using both together.

upvoted 3 times

An administrator has discovered that space utilized by VMs does not decrease after deleting files and folders within the VMs. The administrator needs to be able to reclaim this space.


Which action could the administrator perform to accomplish this task?

- A. Enable Storage I/O Control.
- B. Enable TRIM/UNMAP for the vSAN cluster.
- C. Reboot the VM to recreate the swap file.
- D. Assign a storage policy with thin provisioning.

Suggested Answer: D

Community vote distribution

B (100%)

 **anujspoddar** Highly Voted 5 years, 1 month ago
B

<https://blogs.vmware.com/virtualblocks/2018/09/10/vmware-vsan-6-7u1-storage-reclamation-trim-unmap/>
upvoted 14 times

 **hsezer** Most Recent 3 years, 1 month ago

I also think B should be the answer, but always thinking if it is a tricky question. I found such an article but not completely the same scenario with the question. I want to encourage people to think in a different manner :)

<https://globalvirtualdiary.wordpress.com/2021/02/20/vsan-6-7-reclaim-disk-space-by-converting-thick-provisioned-vms-to-thin-on-vsan/>
upvoted 1 times


 **Lazylinux** 3 years, 4 months ago

Selected Answer: B

Is correct as per my notes
upvoted 1 times

 **Lazylinux** 3 years, 5 months ago

For sure B..No way D as even if you thin provision, once the OS grows it doesnt reclaim space because it is Thin provisioned, actually its current size remains same even if you delete files/directories and Hence B is correct
upvoted 1 times


 **PinPinPoola** 4 years, 3 months ago

Yes, B is correct.

TRIM/UNMAP can also be enabled using PowerCLI.

Query TRIM/UNMAP status: Get-Cluster | Where-Object {\$_.VsanEnabled -eq \$true} | Get-VsanClusterConfiguration | Select-Object Name,GuestTrimUnmap | Format-Table -AutoSize

Enable TRIM/UNMAP: Get-Cluster | Where-Object {\$_.VsanEnabled -eq \$true} | Set-VsanClusterConfiguration -GuestTrimUnmap:\$true
upvoted 2 times

 **diegof1** 4 years, 4 months ago

The correct answer is B

vSAN 6.7 U1 and later support guest operating system initiated ATA Trim and SCSI Unmap requests for data included in the VM virtual disks. The ESXi host issues Trim and Unmap requests to the storage devices for VM swap files and directories mapped to vSAN objects on thin- provisioned virtual disks.

Performance details about Trim/Unmap can be viewed in the Trim/Unmap throughput or latency Performance graphs for each host belonging to a vSAN cluster.

To view these charts, click the Monitor tab for the host and click vSAN > Performance > BACKEND.

The Trim/Unmap feature is not enabled by default. You enable Trim/Unmap using the RVC Console.

Trim/Unmap is supported but disabled by default in vSAN 6.7 U1. The following performance details are available for Trim/Unmap:

- Recovery Unmap Throughput
- Unmap Throughput
- Latency
- Disable

upvoted 3 times

🗨️ 👤 **diegof1** 4 years, 4 months ago

The correct answer is B

vSAN 6.7 U1 and later support guest operating system initiated ATA Trim and SCSI Unmap requests for data included in the VM virtual disks. The ESXi host issues Trim and Unmap requests to the storage devices for VM swap files and directories mapped to vSAN objects on thin- provisioned virtual disks.

Performance details about Trim/Unmap can be viewed in the Trim/Unmap throughput or latency Performance graphs for each host belonging to a vSAN cluster.

To view these charts, click the Monitor tab for the host and click vSAN > Performance > BACKEND.

The Trim/Unmap feature is not enabled by default. You enable Trim/Unmap using the RVC Console.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual

upvoted 2 times

🗨️ 👤 **MIP** 4 years, 9 months ago

I would go with B.

Even if you "assign a storage policy with thin provisioning" (vSAN default policy) the space won't be reclaimed if the TRIM/UNMAP feature is not enabled for the vSAN cluster (even though it should be initially enabled, but it could be disabled).

This feature is enabled initially globally for the cluster:

<https://blogs.vmware.com/virtualblocks/2018/09/10/vmware-vsan-6-7u1-storage-reclamation-trim-unmap>

TRIM/UNMAP is enabled per vSAN cluster using the RVC Console:

<https://storagehub.vmware.com/t/vsan-space-efficiency-technologies/unmap-trim-space-reclamation-on-vsan-1/>

upvoted 3 times

🗨️ 👤 **jasonv** 4 years, 10 months ago

B trim / unmap

upvoted 3 times

🗨️ 👤 **Jackinmany** 4 years, 11 months ago

By default it is thin disk why need to assign storage policy with thin ?

upvoted 2 times

🗨️ 👤 **adelbelkis2** 5 years ago

B is correct

upvoted 3 times

End-users are complaining that an application running on an all-flash vSAN datastore is not performing well. Which performance chart should an administrator view to troubleshoot this issue?

- A. Recovery Write IOPS
- B. Resync Latency
- C. Read Cache Read IOPS
- D. Write Buffer Free Percentage

Suggested Answer: C

Reference:

<https://docs.vmware.com/en/VMware-vSphere/6.7/vsan-673-monitoring-troubleshooting-guide.pdf>

🗳️ 👤 **Jonesythegreat** 3 years, 3 months ago

On a normal basis you should not see resyncs occurring in your vSAN environment. It makes more sense that your write cache is sized too small (even though all flash) and can't destage fast enough to keep up with the application.

upvoted 1 times

🗳️ 👤 **Jonesythegreat** 3 years, 4 months ago

I believe D is the answer. If your caching tier can't drain fast enough you would see performance issues.

upvoted 1 times

🗳️ 👤 **Lazylinux** 3 years, 5 months ago

B is correct

upvoted 1 times

🗳️ 👤 **nsm_t1** 3 years, 7 months ago

B is corrcet answer:

<https://kb.vmware.com/s/article/2150101>

upvoted 1 times

🗳️ 👤 **hsezer** 3 years, 1 month ago

But none of the scenarios are mentoined in the question. It should be defined.

upvoted 1 times

🗳️ 👤 **diegof1** 4 years, 4 months ago

I think B is the correct answer.

upvoted 1 times

🗳️ 👤 **MIP** 4 years, 9 months ago

I think B. I would always look first for latencies (and congestions).

A - high recovery write IOPS show you that something has happened in the cluster and data is recovered to cache or capacity SSDs. big failures could influence the performance...

C - the metric "read cache read IOPS" doesn't exist. I only know the local client cache hit IOPS metric.

D - low write buffer free percentage normally should not influence the performance in an all-flash cluster. colder write buffer data will get destaged from cache to capacity SSDs.

upvoted 1 times

🗳️ 👤 **jasonv** 4 years, 10 months ago

agree with this, B sounds as the most util answer

upvoted 1 times

🗳️ 👤 **LCOJ** 4 years, 11 months ago

I belive that B is the right answer: You can reduce the number of IOPS used to perform resynchronization on disk groups in the vSAN cluster.

Resynchronization throttling is a cluster-wide setting, and it is applied on a per disk group basis.

If VMs are not responding due to latency caused by resynchronization, you can throttle the number of IOPS used for resynchronization. Consider resynchronization throttling only if latencies are rising in the cluster due to resynchronization, or if resynchronization traffic is too high on a host.

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vsan-monitoring.doc/GUID-8D81FCF6-AC9A-4C2C-A8AC-DE50B9965054.html>

upvoted 3 times

Which two features are available with the Cluster Quickstart wizard? (Choose two.)

- A. Deploy a vSAN Witness appliance
- B. Update storage controller drivers
- C. Check for errors and inconsistencies
- D. Setup lockdown mode for vSAN hosts
- E. Perform pre-check evaluation for data migration

Suggested Answer: BC

Reference:

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vsan-planning.doc/GUID-CF9767B6-B3F5-4787-9AF3-D661987AE525.html>

  **Mzoear** Highly Voted 5 years ago

i think C & D is the correct answer
upvoted 6 times

  **manooflon** 4 years, 12 months ago


D . is ok lockdown mode can be configured at step 4 advanced options host options
<https://www.youtube.com/watch?v=5Qc6BYlet3I>
upvoted 5 times

  **diegof1** Most Recent 4 years, 4 months ago

I agree with C & D
upvoted 1 times

  **mariogg85** 4 years, 10 months ago

Agree with C & D
upvoted 1 times

  **LC0J** 4 years, 11 months ago

I agree with the answer C & D as well, they are things that you can do during quick wizard process.
upvoted 2 times

A system administrator running a 6-node cluster is trying to enable the deduplication and compression feature. During the process of switching it on, the General vSAN error occurs.

What is the possible reason the General vSAN error occurred?

- A. Force provisioning setting is required.
- B. The cluster is a hybrid environment.
- C. Storage policies require FTT=1.
- D. The cluster is an all-flash environment.

Suggested Answer: B

Reference:

<http://www.optiodata.com/documents/optio/datasheets/dell-emc-vxrail-appliance-techbook.pdf>

Community vote distribution

B (100%)

🗲️ 👤 **Lazylinux** 3 years, 5 months ago

Selected Answer: B

is correct

upvoted 1 times

🗲️ 👤 **diegof1** 4 years, 4 months ago

B is the correct answer.

Deduplication and compression reduces the amount of data stored on the capacity drives by single-instancing the redundant data stored on each disk group.

Deduplication and compression:

- Are enabled at the cluster level with a single check box
- Have per-cluster settings
- Are available for clusters running:
 - All-flash configurations
 - Advanced licensing
- Are disabled by default

Deduplication and compression is prevented on encrypted VMs.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual

upvoted 3 times

🗲️ 👤 **atechnicznik** 4 years, 7 months ago

Correct

upvoted 2 times

The vSAN host memory needs to be replaced and the entire operation is expected to take 90-120 minutes.

Which vSAN advanced setting can be adjusted to avoid rebuild operations during the host hardware maintenance?

- A. Site Read Locality
- B. Forced Provisioning
- C. Object Repair Timer
- D. Thin Swap

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **LCOJ** Highly Voted 4 years, 11 months ago

I agree with C answer

<https://blogs.vmware.com/virtualblocks/2018/10/29/a-closer-look-at->

[emm/#:~:text=Object%20Repair%20Timer,absent%20objects%20on%20remaining%20hosts.](#)

upvoted 6 times

🗨️ 👤 **Lazylinux** Most Recent 3 years, 4 months ago

Selected Answer: C

Is correct

upvoted 1 times

🗨️ 👤 **Lazylinux** 3 years, 5 months ago

C for sure

upvoted 1 times

🗨️ 👤 **AllenHuang** 4 years, 1 month ago

Agree,C

upvoted 1 times

🗨️ 👤 **diegof1** 4 years, 4 months ago

C is correct answer.

When a host is in maintenance mode (unless the Full data migration option was selected) the objects on the host will be considered "absent." vSAN will wait 60 minutes until initiating a rebuild of the absent objects on remaining hosts. This is because vSAN is not certain if the failure is transient or permanent. If the host will only be in maintenance mode for a few minutes it doesn't make sense to completely rebuild all the objects on a different host. If the host will be offline for more than the default 60 minutes the administrator can modify the vSAN Object repair timer. As of vSAN 6.7 U1 this setting is available in the UI. This is now a cluster-wide setting and will apply to all hosts in the cluster.



upvoted 1 times

You are designing a vSAN stretched cluster solution.

What is the maximum number of nodes possible for this solution?

- A. 15 per data site + 1 witness host
- B. 30 per data site + 1 witness host
- C. 32 per data site + 1 witness host
- D. 64 per data site + 1 witness host



Suggested Answer: A

  **scotland81** Highly Voted 5 years, 1 month ago

Correct answer is A, in a stretched cluster: <https://storagehub.vmware.com/t/vmware-vsan-6-7-technical-overview/stretched-clusters-30/>
upvoted 16 times

  **Lazylinux** Most Recent 3 years, 5 months ago

A for sure and as of vSAN 7 it is $20+20+1 = 41$
upvoted 2 times

  **diegof1** 4 years, 4 months ago

A is correct.

Stretched clusters consist of two active data sites and one witness site. Each site is its own fault domain supporting a fault tolerance of 1 across the active data sites.

A stretched cluster can have up to 15 hosts at each site and must be in communication with a witness host, which brings the host maximum to $15 + 15 + 1$ witness.



Each site should contain the same number of hosts.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual

upvoted 3 times

  **Lance_D** 4 years, 6 months ago

Additional link supporting A: $15+15+1$ as the correct option/answer
upvoted 2 times

  **Lance_D** 4 years, 6 months ago

Answer is A is correct. Please see the link & note below:

[https://download3.vmware.com/vcat/vmw-vcloud-architecture-toolkit-spv1-](https://download3.vmware.com/vcat/vmw-vcloud-architecture-toolkit-spv1-webworks/index.html#page/Storage%20and%20Availability/Architecting%20VMware%20vSAN%206.2/Architecting%20Virtual%20SAN%206.2.2.101.html)

[webworks/index.html#page/Storage%20and%20Availability/Architecting%20VMware%20vSAN%206.2/Architecting%20Virtual%20SAN%206.2.2.101.html](https://download3.vmware.com/vcat/vmw-vcloud-architecture-toolkit-spv1-webworks/index.html#page/Storage%20and%20Availability/Architecting%20VMware%20vSAN%206.2/Architecting%20Virtual%20SAN%206.2.2.101.html)

Note:

With vSAN 6.2, stretched clusters have been enhanced to simplify the creation of the configuration. A new graphical configuration wizard assists with the configuration as appropriate. vSAN stretched clustering is a specific configuration implemented in environments where disaster/downtime avoidance is a key requirement. However, the maximum number of hosts in a stretch cluster configuration remains at 31, where Site 1 contains 15 hosts, Site 2 contains 15 hosts, and site 3 contains the witness host or virtual appliances



upvoted 1 times

  **sugisho** 4 years, 9 months ago

i believe 30 hosts

<https://storagehub.vmware.com/t/vmware-vsan-6-7-technical-overview/stretched-clusters-30/>

upvoted 1 times

  **sugisho** 4 years, 9 months ago

mistake

A maximum of thirty (30) hosts may be used in a single vSAN Stretched Cluster across the data sites. In vSAN versions up to 6.5, this was fifteen (15) per site. With the introduction of Site Affinity rules that places data on only one data site or the other, it is possible to have a vSAN Stretched Cluster deployment that does not have an equal number of hosts per site.

upvoted 1 times

🗨️ 👤 **alsmk2** 4 years, 9 months ago

B

Have done a couple of 30+1 stretched deployments - 15 hosts per site as others have already stated.

Question doesn't state specifically the maximum nodes for 1 site, so for the entire solution, B is correct.

upvoted 1 times

🗨️ 👤 **alsmk2** 4 years, 8 months ago

A is the correct answer actually - misread the question 15 nodes per data site + 1. Two data sites = 30+1

upvoted 1 times

🗨️ 👤 **MIP** 4 years, 9 months ago

A

In vSAN Stretched Cluster maximum number of data hosts on each site is 15.

Limit: 30 (15 on each site x 2)

<https://configmax.vmware.com/guest?vmwareproduct=vSphere&release=vSphere%206.7&categories=7-0>

upvoted 1 times

🗨️ 👤 **oklanmayor** 4 years, 10 months ago

A, 15 per site + 1 witness

upvoted 2 times

🗨️ 👤 **avilis** 4 years, 11 months ago

B. 30 it is.

upvoted 1 times

🗨️ 👤 **Kywik** 4 years, 11 months ago

Correct answer is B.

A maximum of thirty (30) hosts may be used in a single vSAN Stretched Cluster across the data sites.

upvoted 1 times

🗨️ 👤 **Kywik** 4 years, 11 months ago

Oups ! It's 30 hosts across all data sites. So correct answer is A (15 nodes per data site).

upvoted 2 times

What are two recommended reasons for configuring a cluster with at least one node in addition to the minimum required number? (Choose two.)

- A. To ensure object accessibility
- B. To support the use of RAID-6
- C. To provide more flexible storage policy options
- D. To provide full protection during maintenance mode operations
- E. To support data at rest encryption on vSAN hybrid clusters

Suggested Answer: CD

Reference:

<https://blogs.vmware.com/virtualblocks/2018/05/24/vsan-deployment-considerations/>

Community vote distribution

CD (100%)

🗳️ **xurooj** 3 years, 2 months ago

C is fine but certainly not D, as all the objects are fully protected with the policy either accessibility or data migrate, the only reason is to rebuild the objects after failure.

upvoted 1 times

🗳️ **Jonesythegreat** 3 years, 4 months ago

This should be A and D. You are not increasing your storage policy options by doing N+1, but you are increasing your availability to recover with failures.

upvoted 1 times

🗳️ **Lundu1995** 3 years, 4 months ago

mhn i'm not really 100% sure if its c&d . it could be a & d as well cause more nodes = more votes on the object and if you have more votes on an object the object is accessible if there is a failure

upvoted 2 times

🗳️ **Lazylinux** 3 years, 4 months ago

Selected Answer: CD

is correct

upvoted 1 times

🗳️ **Lazylinux** 3 years, 5 months ago

C and D for sure

upvoted 1 times

🗳️ **AllenHuang** 4 years, 1 month ago

Agree,C&D

upvoted 1 times

🗳️ **diegof1** 4 years, 4 months ago

C & D are the correct answers

upvoted 2 times

🗳️ **atechnicznik** 4 years, 7 months ago

I agree with the answers

upvoted 1 times

vCenter reports a number of vSAN network alarms on a 2-node vSAN Direct Connect cluster. The test pings show irregularities in the ping results. Using vSAN Health service, which two checks should be used to identify the network root cause? (Choose two.)

- A. Physical disk operation health
- B. Network latency
- C. vSAN extended configuration
- D. vSAN disk balance
- E. vSAN basic unicast connectivity

Suggested Answer: AB

Reference:

<https://vsan-essentials.gitbooks.io/vsan-6-2/content/chapter10.html>

Community vote distribution

BE (100%)

  **Mzoear** Highly Voted 5 years ago

i think B & E is the correct answer



upvoted 7 times

  **Lazylinux** Most Recent 3 years, 5 months ago

Selected Answer: BE



Are correct

upvoted 1 times

  **diegof1** 4 years, 4 months ago

B & E are the correct answer.

upvoted 2 times

  **wtkc** 4 years, 9 months ago

E is only pinging. So we can not detect root cause.



Probably A and B are correct.

if physical disk is dying, maybe ping is irregular. So I choose A.

The other network health check results should be examined to narrow down the root cause of the misconfiguration.

<https://storagehub.vmware.com/t/vmware-r-vsan-tm-network-design/vmotion-basic-unicast-connectivity-check/>

upvoted 1 times

  **alsmk2** 4 years, 8 months ago



A physical disk dying would be storage, not network. B & E are correct.

upvoted 1 times

  **LCOJ** 4 years, 11 months ago

I agree with B & E are correct.

upvoted 4 times

  **Ady_14** 4 years, 11 months ago

B & E are correct.

upvoted 3 times

A vSAN administrator is designing a new all-flash vSAN cluster. The cluster will host read intensive applications. Which factor should be included in the design to improve read performance?

- A. Multiple disk groups
- B. Large capacity disks
- C. Large cache disks
- D. Large controller cache

Suggested Answer: B

Reference:

<https://vsan-essentials.gitbooks.io/vsan-6-2/content/chapter4.html>

Community vote distribution

A (100%)

🗳️ 👤 **Darius_Th3D0G** Highly Voted 5 years, 1 month ago

I think the correct answer is A: "multiple disk groups". <https://blogs.vmware.com/virtualblocks/2019/04/18/vsan-disk-groups/>
upvoted 7 times

🗳️ 👤 **Lazylinux** Most Recent 3 years, 4 months ago

Selected Answer: A

is correct

upvoted 1 times

🗳️ 👤 **Lazylinux** 3 years, 5 months ago

A for sure

upvoted 1 times

🗳️ 👤 **diegof1** 4 years, 4 months ago

A is the correct answer.

Breaking up the disk groups has the following advantages:

- Redundancy: In example one, a single disk groups can result in a host failure.
- Flexibility: Scale up or scale out are both viable options.
- Performance: More cache devices means more I/O paths.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - Design: One or Multiple Disk Groups Per Host Section

upvoted 2 times

🗳️ 👤 **wtkc** 4 years, 9 months ago

In all flash mode, there is no read cache and directly read to capacity disks.

So, Adding or Splitting cache don't improve read performance.

I choose B.

upvoted 2 times

🗳️ 👤 **alsmk2** 4 years, 8 months ago

Adding multiple disk groups = more IO paths = more performance. A is correct.

upvoted 3 times

🗳️ 👤 **LCOJ** 4 years, 11 months ago

A agree thay A is the correct answer

upvoted 4 times

🗳️ 👤 **Ady_14** 4 years, 11 months ago

A is correct.

upvoted 3 times

Which storage policy action results in a component resync?

- A. Enabling object checksum
- B. Changing object space reservation to thick provisioning
- C. Adding an IOPS Limit rule to a storage policy
- D. Changing the failure tolerance method

Suggested Answer: D

Community vote distribution

D (100%)

  **LCOJ** Highly Voted 4 years, 11 months ago

The answer D is correct:

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID-04ED9C07-BC61-4E3C-9188-9D2E460C9CA3.html>
upvoted 7 times

  **Lazylinux** Most Recent 3 years, 4 months ago

Selected Answer: D



is correct

upvoted 1 times

  **Lazylinux** 3 years, 5 months ago

D for sure

upvoted 1 times

  **diegof1** 4 years, 4 months ago

D is the correct answer.

vSAN initiates resynchronization either when a hardware device, host, or network fails, or a host is placed into maintenance mode in the vSAN cluster.



Resynchronization triggers:

- Editing a VM storage policy
- Restarting a host after a failure
- Recovering hosts from a permanent or long-term failure (60 minutes by default)
- Evacuating data by using the Full data migration mode before placing a host in maintenance mode
- Exceeding the capacity threshold of a capacity device (threshold level of 80%)

So changing the Failure Tolerance Method is to change the VM storage policy.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - Resynchronizing Components Section

upvoted 1 times

  **diegof1** 4 years, 4 months ago

D is the correct answer.



vSAN initiates resynchronization either when a hardware device, host, or network fails, or a host is placed into maintenance mode in the vSAN cluster.

Resynchronization triggers:

- Editing a VM storage policy
- Restarting a host after a failure
- Recovering hosts from a permanent or long-term failure (60 minutes by default)
- Evacuating data by using the Full data migration mode before placing a host in maintenance mode
- Exceeding the capacity threshold of a capacity device (threshold level of 80%)



Taken from vSAN 6.7 Deploy and Manage Lecture Manual - Resynchronizing Components Section

upvoted 1 times

  **alsmk2** 4 years, 9 months ago

D is the most obvious, but A and B are both listed in various places. Bad question.

upvoted 4 times

  **MIP** 4 years, 9 months ago

I'll go with D (but it could also be B)

VMworld 2019 Session "HCI3450BE vSAN best practices"

=> Which policy changes require a rebuild? (Slide 43)

- increasing/decreasing stripe width
- changing RAID protection aka FTM (RAID1-to/from-5/6, RAID5-to/from-6)
- increasing the object space reservation (OSR can only be 0% or 100% when deduplication is enabled)

upvoted 3 times

  **mariogg85** 4 years, 10 months ago

I would say the right answer could be A

<https://cormachogan.com/2018/02/21/policy-changes-can-trigger-rebuild-vsan/>

Enabling Object Checksum results on a component resync

upvoted 1 times

vSAN is configured with a storage policy with a Failures To Tolerate (FTT) set to three.
Which type of redundancy is being used?

- A. RAID-0 (Striping)
- B. RAID-1 (Mirroring)
- C. RAID-6 (Erasure Coding)
- D. RAID-5 (Erasure Coding)

Suggested Answer: B

Reference:

https://www.sbarjatiya.com/notes_wiki/index.php/About_VSAN_6.6

Community vote distribution

B (100%)

🗨️ 👤 **LCOJ** Highly Voted 4 years, 11 months ago

The answer is right, B:

<https://storagehub.vmware.com/t/vsan-space-efficiency-technologies/host-requirements-1/>

upvoted 6 times

🗨️ 👤 **bhaitraining** 4 years, 5 months ago

True. FTT=3 cannot be used with erasure coding and mirroring (RAID1) must be used in case of FTT=3.

upvoted 1 times

🗨️ 👤 **RRK06** Most Recent 3 years, 1 month ago

Selected Answer: B

B only supports FTT =3

upvoted 1 times

🗨️ 👤 **Lazylinux** 3 years, 5 months ago

B for sure

upvoted 1 times

🗨️ 👤 **diegof1** 4 years, 4 months ago

B is the correct answer.

RAID 5/6 (erasure coding) does not support a failure tolerance of 3.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - RAID 1 Mirroring Versus RAID 5/6 Erasure Coding Section

upvoted 1 times

Which default VMware vSphere High Availability cluster setting must be changed when using a vSAN stretched cluster?

- A. Datastore with All Paths Down
- B. Datastore with Permanent Device Loss
- C. Response for Host Isolation
- D. Enable Host Monitoring



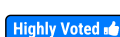
Suggested Answer: D

Reference:



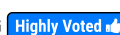
<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID-D68890D8-841A-4BD1-ACA1-DA3D25B6A37A.html>

Community vote distribution

C (100%)

  **LFC**  5 years ago

Answer is C - Enable host monitoring is turned on by default. For a stretched cluster, host isolation needs to be changed from default to Power off and restart VMs. <https://storagehub.vmware.com/t/vsan-stretched-cluster-guide/cluster-settings-vsphere-ha-5/>
upvoted 11 times

  **Darius_Th3D0G**  5 years, 1 month ago

I think the right answer is C.
<http://www.vexpertconsultancy.com/2018/03/vmware-vsan-vsphere-ha-best-practices/>
upvoted 8 times

  **khannat**  2 years, 8 months ago

The answer is A : Note that all vSAN clusters participating in HCI Mesh (a feature introduced in vSAN 7 U1 and enhanced in vSAN 7 U2) must also configure the "All Paths Down (APD) Failure Response" found in "HA Cluster Settings" in vCenter Server, and ensure it is enabled and configured for proper APD
upvoted 1 times

  **Lazylinux** 3 years, 4 months ago



Selected Answer: C

is correct
upvoted 1 times

  **Lazylinux** 3 years, 5 months ago

Selected Answer: C

For sure C....Enable host monitoring is turned on by default. For a stretched cluster, host isolation needs to be changed from default to Power off and restart VMs.
upvoted 1 times



  **dikboy** 3 years, 7 months ago

I think C refer to this doc

"When a vSAN node becomes isolated, vSAN will power off virtual machines, but will not restart them. vSphere HA is used to restart these virtual machines on hosts that have not been isolated. "

<https://core.vmware.com/resource/vsan-stretched-cluster-guide#sec7300-sub2>

upvoted 1 times

  **diegof1** 4 years, 4 months ago

C is the correct answer.

vSphere HA Setting: Host Isolation Response (Default: Remain Powered On)

Recommended Value for vSAN: Power off and restart VMs

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - vSphere HA Settings with vSAN Section

upvoted 1 times

🗨️ 👤 **Ganga_red** 4 years, 9 months ago

The answer is D. Host Monitoring is by default is disabled. Post enabling the Host monitoring we will get a option of Host isolation response.

<https://storagehub.vmware.com/t/vsan-stretched-cluster-guide/host-monitoring-5/>

<https://cormachogan.com/2015/09/16/vsphere-ha-settings-for-vsan-stretched-cluster/#:~:text=Host%20Monitoring&text=This%20will%20make%20your%20virtual,the%20cluster%2C%20and%20are%20healthy.>

upvoted 4 times

🗨️ 👤 **LookingUp** 4 years, 7 months ago

By default for a new Cluster Enable Host Monitoring is Enabled, only VMware HA is disabled. You would need to change the "Response for Host Isolation" setting making C the right answer.

upvoted 2 times

🗨️ 👤 **DenZn** 4 years, 11 months ago

Agree with answer C

upvoted 4 times

🗨️ 👤 **Ady_14** 4 years, 11 months ago

C is the correct answer.

upvoted 3 times

🗨️ 👤 **adelbelkis2** 5 years ago

I believe it's C

upvoted 3 times

🗨️ 👤 **Darius_Th3D0G** 5 years, 1 month ago

Sorry for my previous message. D is the correct answer.

upvoted 2 times

The cluster level backend IOPS performance graph shows a higher-than-average number of IOPS.

What is a possible reason for this?

- A. Data resynchronizations
- B. Low volume of vSAN vital memory pools
- C. DRS invoking multiple vMotion migrations
- D. Destaging congestion

Suggested Answer: A

Reference:

<https://kb.vmware.com/s/article/2144493>

  **Lazylinux** 3 years, 5 months ago

A for sure...

vSAN cluster read IOPS of resync traffic, including policy change, repair, maintenance mode / evacuation and rebalance from resyncing objects in the perspective of vSAN backend.

<https://kb.vmware.com/s/article/2144493#vSANBackendGraph>

upvoted 1 times



  **Dzma** 4 years, 3 months ago

A for sure!

Measuring at the cluster level provides context, and helps identify influencing factors. Remember that vSAN is a cluster based storage solution, and VM data is not necessarily always residing on the host that the VM is residing. While it might seem odd to check look at performance metrics at this level as the second step, it can often help provide an understanding to the level of activity across the cluster. For example, perhaps a VM latency spikes occur during the middle of the night. After identifying the VM level statistics, viewing the cluster level statistics might show a substantial amount of noise coming from other VMs, or perhaps backend resynchronization traffic. To view the vSAN cluster based performance metrics, highlight the cluster, click on Monitor > vSAN > Performance to view the respective VM, backend, or iSCSI performance metrics.

<https://core.vmware.com/resource/troubleshooting-vsan-performance>

upvoted 1 times

  **diegof1** 4 years, 4 months ago

A is right.

upvoted 1 times

An administrator of a vSAN 6.5 environment is planning a path to move to unicast for the vSAN network.
How will the administrator achieve this?

- A. Enable LACP on vSAN switch ports.
- B. Upgrade vCenter and vSAN nodes to vSAN 6.6 or later.
- C. Enable unicast on the vSAN interfaces.
- D. Disable multicast for vSAN using ESXCLI commands.

Suggested Answer: B

Community vote distribution

B (100%)

  **LCOJ**  4 years, 11 months ago

The answer B is right, as of version 6.6 vSAN does not use Multicast anymore. It has been updated to work with Unicast by default.

<https://blogs.vmware.com/virtualblocks/2017/04/11/goodbye-multicast/>



upvoted 6 times

  **Lazylinux**  3 years, 5 months ago

Selected Answer: B

Is correct

upvoted 1 times

  **diegof1** 4 years, 4 months ago

B is the correct answer.

Unicast in vSAN network: In vSAN 6.6 and later releases, multicast is not required on the physical switches that support the vSAN cluster. You can design a simple unicast network for vSAN.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - vSAN Network Features section

upvoted 1 times

When using vSAN Encryption, where does the Key Encryption Key persistently reside?

- A. in /etc/vmware/ssl on each vSAN host
- B. in the KMS server
- C. in a VM configuration file on vSAN
- D. in the vCenter Server cache


Suggested Answer: C

Reference:

<https://blogs.vmware.com/virtualblocks/2018/07/13/understanding-ve-booting-w-vc-unavailable/>

Community vote distribution

B (100%)

 **Darius_Th3D0G** Highly Voted 5 years, 1 month ago

I think B is the correct anaswer:

<https://blogs.vmware.com/virtualblocks/2018/07/13/understanding-ve-booting-w-vc-unavailable/>

upvoted 7 times

 **LC0J** Highly Voted 4 years, 11 months ago

I agree with B as per state below

vCenter Server requests an AES-256 Key Encryption Key (KEK) from the KMS. vCenter Server stores only the ID of the KEK, but not the key itself.

Each ESXi host uses the KEK to encrypt its DEKs, and stores the encrypted DEKs on disk. The host does not store the KEK on disk. If a host reboots, it requests the KEK with the corresponding ID from the KMS. The host can then decrypt its DEKs as needed.


upvoted 7 times

 **RRK06** Most Recent 3 years, 1 month ago

Selected Answer: B

Key is always with Key server, it will be retrieved during need by ESXi host and vcenter

upvoted 1 times

 **Lundu1995** 3 years, 4 months ago


B is correct!

upvoted 1 times

 **Lazylinux** 3 years, 5 months ago

B for sure

upvoted 1 times

 **diegof1** 4 years, 4 months ago


B is the correct answer.

When encryption is enabled on a vSAN cluster:

1. vCenter Server requests an AES-256 KEK from the KMS. vCenter Server stores only the ID of the KEK.
2. vCenter Server sends the KEK ID to all hosts.
3. Hosts use the KEK ID to request the KEK from the KMS.
4. Hosts create a unique DEK for each drive.
5. The vSAN datastore is encrypted with each drive having its own DEK.
6. KMS generates a single Host Key HEK sent to all hosts in the cluster used for encrypting core dumps.

Taken from vSAN 6.7 Deploy and Manage - Encryption Key Generation section

upvoted 3 times

 **Ady_14** 4 years, 11 months ago

B is the right Answer

upvoted 2 times

 **adelbelkis2** 5 years ago

I believe B is correct from the link below "The KEK and Host Key are placed in memory in the key cache. These keys are not persistently stored on the vSAN hosts"

upvoted 2 times

What is a mandatory requirement for a vSAN stretched cluster?

- A. Two witness hosts protected by vSphere Fault Tolerance
- B. Minimum of two guaranteed lines of 10Gbps bandwidth
- C. vSphere Replication configured between the data sites
- D. Latency between data sites must be less than or equal to 5ms

Suggested Answer: D

Reference:

<https://vsan-essentials.gitbooks.io/vsan-6-2/content/chapter8.html>

Community vote distribution



D (100%)

  **LCOJ** Highly Voted 4 years, 11 months ago

The answer D is correct:

<https://storagehub.vmware.com/t/vmware-vsan-6-7-technical-overview/stretched-clusters-30/>

upvoted 5 times

  **Lazylinux** Most Recent 3 years, 5 months ago

Selected Answer: D



Because it is

upvoted 1 times

  **AllenHuang** 4 years, 1 month ago

D is correct

upvoted 1 times

  **diegof1** 4 years, 4 months ago

D is the correct answer.

The data sites should have very low (≤ 5 ms) latency (RTT) and high bandwidth. Witness requirements are relaxed.

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - Single-Site Versus Stretched Clusters section

upvoted 1 times

An architect is asked to design a vSAN cluster with these requirements:

- Supports erasure coding storage policy
- Minimized inaccessibility from single device failures
- Ability to re-protect data after a failure

The desired configuration should be able to survive a single failure.

Which configuration satisfies the design requirements?

- A. 3-node hybrid vSAN cluster
- B. 4-node hybrid vSAN cluster
- C. 4-node all-flash vSAN cluster
- D. 5-node all-flash vSAN cluster

Suggested Answer: D

Community vote distribution

D (100%)

  **LFC** Highly Voted 5 years ago

Answer D - Min no of nodes in Raid 5 _ erasure coding is 4 , To satisfy ability to re-protect data after a failure Need one additional node. So the total node count is 5.

upvoted 7 times

  **Lazylinux** Most Recent 3 years, 5 months ago

Selected Answer: D



D for sure

upvoted 1 times

  **AllenHuang** 4 years, 1 month ago

D is correct



upvoted 1 times

  **diegof1** 4 years, 4 months ago

D is the correct answer

RAID 5 just needs 4 hosts to support FTT=1. However, the statement said -Ability to re-protect data after a failure, to support this you have to ensure 4 hosts available after one failure. Thus, you will need 5 hosts in the cluster.

upvoted 3 times

  **LCOJ** 4 years, 11 months ago

I agree with D answer.

upvoted 3 times

  **adelbelkis2** 5 years ago

I believe it's D.

Because of : Ability to re-protect data after a failure, if we have 4 nodes we won't be able to reprotect data as we'll have less than the minimim hosts which is "4" (Raid 5)

upvoted 3 times

A new vSAN cluster needs to run 900 VMs and protect against one failure.
What is the minimum number of nodes needed to meet the requirement?

- A. 3
- B. 4
- C. 5
- D. 6

Suggested Answer: A

Community vote distribution

D (100%)

🗳️ 👤 **Mohamed_Atef** Highly Voted 4 years, 11 months ago

Correct Answer is D! As a vSAN maximum, each node can host up to 200 VM which means 5 nodes are required to host the 900 VMs, and to tolerate 1 node failure the total number of nodes should be 6
upvoted 14 times

🗳️ 👤 **DenZn** 4 years, 11 months ago

True! Thank you for it!
upvoted 2 times

🗳️ 👤 **guruto** 4 years, 5 months ago

I agree.
<https://configmax.vmware.com/guest?vmwareproduct=vSphere&release=vSphere%206.7&categories=7-0>
upvoted 1 times

🗳️ 👤 **Venkyrag** 4 years, 5 months ago

I agree
upvoted 1 times

🗳️ 👤 **LC0J** Highly Voted 4 years, 11 months ago

I agree with D answer: as per configuration maximum each vSAN node can host up to 200 VM, this way required 6 nodes to be self resilient.
<https://configmax.vmware.com/guest?vmwareproduct=vSphere&release=vSphere 6.7&categories=7-0>
upvoted 5 times

🗳️ 👤 **Lazylinux** Most Recent 3 years, 5 months ago

Selected Answer: D

Is correct
upvoted 1 times

🗳️ 👤 **Lazylinux** 3 years, 5 months ago

Correct Answer D
upvoted 1 times

🗳️ 👤 **AllenHuang** 4 years, 1 month ago

D is correct
upvoted 1 times

🗳️ 👤 **diegof1** 4 years, 4 months ago

The correct answer is D

vSAN supports up to 200 virtual machines per host. So you need to ensure 6 hosts to support 1200. Now if FTT=1 the rest 5 host should be able to support 1000 VMs (covering the 900 VMs required)
upvoted 1 times

🗳️ 👤 **DenZn** 4 years, 11 months ago

Answer A is right:
Raid 1 - FTT=1 - Minimum number of Hosts = 3

<https://blogs.vmware.com/virtualblocks/2018/05/24/vsan-deployment-considerations/>

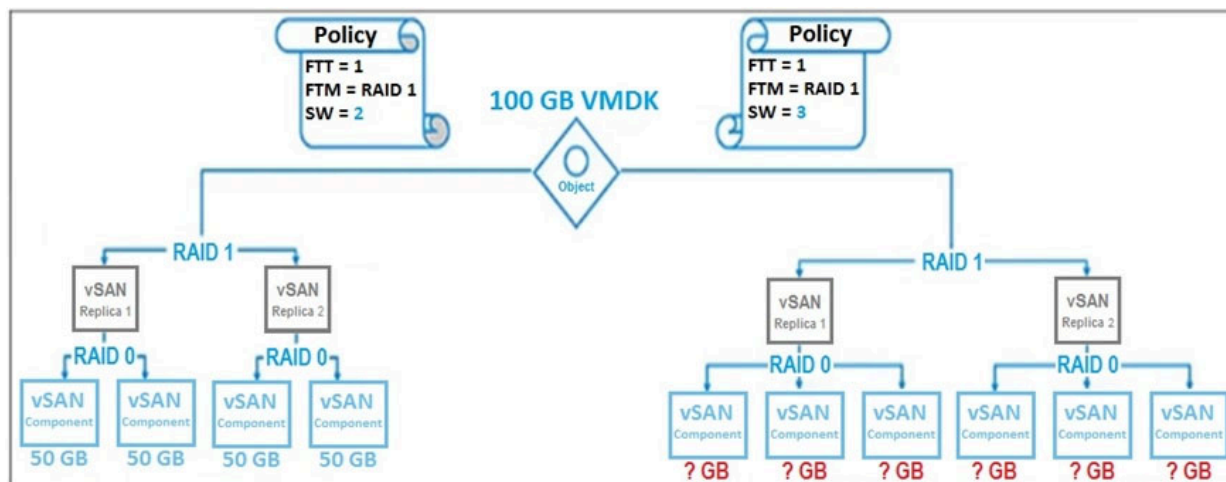
upvoted 1 times

  **Mimmo2020** 4 years, 12 months ago

Correct Answer should be: D

upvoted 2 times

Refer to the exhibit.



In a 2-node vSAN environment, a storage administrator has setup a vSAN storage policy.

When the stripe width is changed to three, what is the approximate component size of each replica marked in red?

- A. 33GB
- B. 50GB
- C. 66GB
- D. 100GB

Suggested Answer: A

Community vote distribution

A (100%)

Lazylinux 3 years, 4 months ago

Selected Answer: A

Is correct

upvoted 1 times

Lazylinux 3 years, 5 months ago

A for sure

upvoted 1 times

AllenHuang 4 years, 1 month ago

A is correct

upvoted 1 times

diegof1 4 years, 4 months ago

A is the correct answer.

100 GB is the size of VMDK, then if you stripe in three different components, each one should have 33,33 GB \approx 33 GB.

upvoted 1 times

diegof1 4 years, 4 months ago

A is the correct answer.

100 GB is the size, then if you strip in tree different components, each one should have 33,33 GB \approx 33 GB.

upvoted 2 times

missing_dll 4 years, 7 months ago

A is correct,

50+50 \approx 33+33+33

upvoted 2 times

Which two factors should be considered when planning usable vSAN datastore capacity? (Choose two.)

- A. Flash Read Cache Reservation percentage
- B. Failure tolerance method
- C. Stripe width
- D. IOPS limit for object
- E. Failures To Tolerate

Suggested Answer: *BE*

Reference:

<https://docs.vmware.com/en/VMware-vSphere/6.7/vsan-671-planning-deployment-guide.pdf>

(23)

Community vote distribution



BE (100%)

  **Lazylinux** 3 years, 5 months ago

Selected Answer: BE

For sure

upvoted 1 times

  **diegof1** 4 years, 4 months ago

B & E are correct.

upvoted 1 times

  **sheriefmb** 4 years, 5 months ago

Agree with the answer

upvoted 3 times

An ESXi host is being re-purposed from a 4-node cluster. The default policy is in use and storage utilization is low. Which maintenance mode should be used to avoid storage policy non-compliance?

- A. Ensure accessibility
- B. No data migration
- C. Proactive rebalance
- D. Full data migration

Suggested Answer: A

Reference:

<https://storagehub.vmware.com/t/vmware-vsan/vsan-2-node-guide/maintenance-mode-consideration-9/>

Community vote distribution

D (100%)

🗳️ 👤 **xurooj** 3 years, 2 months ago

ensure accessibility will move only that data which is not on any other disks/groups.

upvoted 1 times

🗳️ 👤 **Lazylinux** 3 years, 4 months ago

Selected Answer: D

is correct

upvoted 1 times

🗳️ 👤 **Lazylinux** 3 years, 5 months ago

D for sure

upvoted 1 times

🗳️ 👤 **diegof1** 4 years, 4 months ago

D is the correct answer.

Full data migration ensures that storage policy won't become non-compliance.

upvoted 4 times

🗳️ 👤 **LCOJ** 4 years, 11 months ago

I agree with the answer D. The host is going to be removed permanently.

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/GUID-521EA4BC-E411-47D4-899A-5E0264469866.html>

upvoted 3 times

🗳️ 👤 **Ady_14** 4 years, 11 months ago

D is the right Answer.

upvoted 3 times

🗳️ 👤 **adelbelkis2** 5 years ago

I believe D.

From the same reference "Ensure accessibility – The preferred option for two-host or three-host vSAN clusters using the default storage policy.

Ensure accessibility guarantees the enough components of the vSAN object are available for the object to remain available. Though still accessible, vSAN objects on two-host or three-host clusters are no longer policy compliant"

upvoted 2 times

🗳️ 👤 **LFC** 5 years ago

Correct answer is D.

The host is going to remove from the cluster permanently.

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/GUID-521EA4BC-E411-47D4-899A-5E0264469866.html>

upvoted 4 times

What would cause a vSAN: Basic (unicast) connectivity check failure in vSAN Health Check?

- A. One or more hosts in the cluster are unable to ping the vSAN VMkernel adapters of other hosts in the cluster.
- B. The throughput of the physical network supporting vSAN traffic is less than the recommended minimum throughput.
- C. The vSAN VMkernel adapter is shared with another service such as vMotion.
- D. vSAN is configured to use multicast to transmit data.

Suggested Answer: A

Reference:

<https://communities.vmware.com/thread/576629>

Community vote distribution

A (100%)

🗄️ 👤 **LCOJ** Highly Voted 👍 4 years, 11 months ago

The answer A is right.

<https://storagehub.vmware.com/t/vmware-r-vsan-tm-network-design/vsan-basic-unicast-connectivity-check/>

upvoted 6 times

🗄️ 👤 **Lazylinux** Most Recent 🕒 3 years, 5 months ago

Selected Answer: A

is correct

upvoted 1 times

🗄️ 👤 **Lazylinux** 3 years, 5 months ago

A for sure

upvoted 1 times

What is the maximum number of failures an object can tolerate with RAID-6?

- A. 2
- B. 3
- C. 4
- D. 6

Suggested Answer: A

Reference:

<https://vsan-essentials.gitbooks.io/vsan-6-2/content/chapter6.html>

Community vote distribution

A (100%)

🗲️ 👤 **Lazylinux** 3 years, 4 months ago

Selected Answer: A

is correct

upvoted 1 times

🗲️ 👤 **Lazylinux** 3 years, 5 months ago

A for sure

upvoted 1 times

🗲️ 👤 **diegof1** 4 years, 4 months ago

A is correct

RAID5/6 support up to FTT=2.

upvoted 1 times

🗲️ 👤 **diegof1** 4 years, 4 months ago

A is correct.

RAID5/6 support just FTT=2.

upvoted 1 times

What is a supported use of the vSAN iSCSI target service?

- A. To provision storage as a VM Raw Device Mapping
- B. To provision storage for third-party hypervisors
- C. To provision storage for an ESXi host
- D. To provision storage for an Oracle RAC cluster

Suggested Answer: D

Reference:

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/GUID-13ADF2FC-9664-448B-A9F3-31059E8FC80E.html>

Community vote distribution

D (100%)

🗳️ 👤 **Lazylinux** 3 years, 5 months ago

Selected Answer: D

For sure

upvoted 1 times

🗳️ 👤 **diegof1** 4 years, 4 months ago

D is right.

The iSCSI target service enables hosts and physical workloads that reside outside the vSAN cluster to access the vSAN datastore. This feature enables an iSCSI initiator on a remote host to transport block-level data to an iSCSI target on a storage device in the vSAN cluster, allowing the discovery of vSAN iSCSI targets from a remote host. vSAN 6.7 and later releases support Windows Server Failover Clustering (WSFC), allowing WSFC nodes to access vSAN iSCSI targets.

To ensure high availability of the vSAN iSCSI target, configure multipath support for your iSCSI application. You can use the IP addresses of two or more hosts to configure the multipath.

Note: The vSAN iSCSI target service does not support other vSphere, ESXi clients, initiators, third-party hypervisors, or migrations using raw device mapping (RDM).

Taken from vSAN 6.7 Deploy and Manage Lecture Manual - vSAN iSCSI Target Service Section

upvoted 2 times

🗳️ 👤 **LCOJ** 4 years, 11 months ago

Answer D is correct, the iSCSI is supported with Oracle RAC, WSFC from Microsoft and to delivery LUNs to physical host out of the box.

<https://storagehub.vmware.com/t/vmware-r-vsan-tm-network-design/iscsi-on-vsan-limitations-and-considerations/>

upvoted 3 times