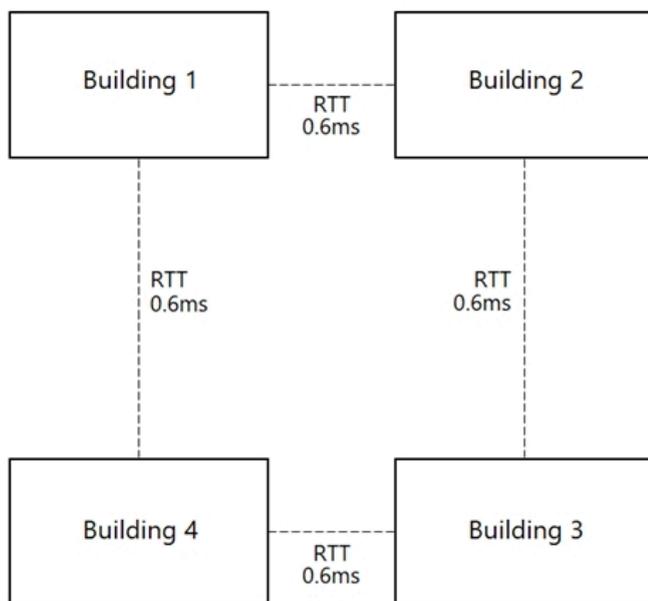




- Expert Verified, Online, **Free**.

Reference the exhibit.

----- 40Gbps, Layer 3 network



A university campus has a site topology illustrated in the exhibit. Adjacent connected buildings are connected with a single 40gb layer 3 network with a RTT of

0.6ms as depicted by each line. An architect wants to configure a fault domain in each building.

What would the architect need change in order to implement a fully supported vSAN Cluster?

- A. Enable a single Layer 2 network across all buildings for vSAN connectivity.
- B. Decrease the existing RTT between the existing links by 0.1ms or more.
- C. Add an additional link to the existing adjacent links to make it redundant.
- D. Enable IPv6 networks across all buildings for vSAN connectivity.

Suggested Answer: C

Community vote distribution

A (100%)

jsi928 1 year, 1 month ago

B. Site to site must be 5ms or less. Layer 2 is recommended and Layer 3 is supported.

<https://docs.vmware.com/en/VMware-vSphere/8.0/vsan-network-design-guide/GUID-F3401655-6EFA-477B-B072-E8F001B50BCC.html>

upvoted 1 times

Akesh7 1 year, 6 months ago

Selected Answer: A

Layer 2 is the basic network requirement

upvoted 1 times

qazxsx 2 years, 8 months ago

I think B is correct.

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vsan-planning.doc/GUID-AFF133BC-F4B6-4753-815F-20D3D752D898.html>

upvoted 2 times

mparayil 3 years, 4 months ago

I think stretching the Layer2 network will be only the reconfiguration, required, because all Building as 2 Connectivity in a ring topology.

upvoted 2 times

MrVMware 3 years, 5 months ago

i cant see how this a supported config ? fault domains are for same site rack failure , so stretched cluster in which case will be 3 building not 4 ?

upvoted 1 times

  **Spiff** 3 years, 3 months ago

I believe the limit on VxRail 4.x is 3 fault domains. Native vSAN doesn't have that 3 fault domain limit.
upvoted 1 times

  **Rabbit117** 3 years, 5 months ago

I think C is correct.
upvoted 2 times

When designing a vSAN Cluster, which three should be considered when sizing a vSAN datastore? (Choose three.)

- A. Storage Policy or Policies to be used
- B. capacity device queue depth
- C. vSAN Slack Space
- D. cache disks with 10% of the Disk Group
- E. On-Disk Format overhead
- F. endurance level of capacity disk used

Suggested Answer: ABC

Reference:

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

Community vote distribution

ACE (100%)

🗨️ 👤 **Lazylinux** 2 years, 8 months ago

Selected Answer: ACE

is correct

upvoted 1 times

🗨️ 👤 **gkoylu** 3 years, 2 months ago

I think the answer is ACE because Queue Depth is about the performance, not capacity...

upvoted 2 times

🗨️ 👤 **Alefin** 3 years, 6 months ago

A-C-E are Correct

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vsan-planning.doc/GUID-581D2D5C-A88F-4318-A8B3-5A5F343F1247.html>

upvoted 2 times

🗨️ 👤 **dfdd1** 3 years, 3 months ago

Why you should consider "endurance level of capacity disk used" when sizing it does make no sense?! I think it should be A-C-D

upvoted 1 times

🗨️ 👤 **Rabbit117** 3 years, 6 months ago

I think A, C and E are correct as these are related to the usable capacity of the storage.

upvoted 2 times

Which must be configured by a Kubernetes user allow for dynamic provisioning of container volumes in vSAN using storage policies?

- A. ReplicaSet
- B. Pod
- C. StorageClass
- D. ReplicationController

Suggested Answer: *C*

Reference:

<https://vmware.github.io/vsphere-storage-for-kubernetes/documentation/policy-based-mgmt.html>

🗨️ 👤 **Alefin** 3 years, 6 months ago

C is correct

<https://vmware.github.io/vsphere-storage-for-kubernetes/documentation/policy-based-mgmt.html>

upvoted 2 times

🗨️ 👤 **Rabbit117** 3 years, 6 months ago

I think C is correct.

<https://cormachogan.com/2019/06/04/kubernetes-storage-on-vsphere-101-storageclass/?amp=1>

upvoted 3 times

An All-Flash vSAN cluster has 4 nodes with this disk group composition on each host:

- ⇒ 1 Å– 800 GB SAS SSD
- ⇒ 6 Å– 3.84 TB SATA SSD

An administrator observes that in an All-Flash cluster, the Write Buffer Free Percentage is consistently low. As a result, there is increased latency experienced by the workloads.

Which remediation steps will resolve this issue?

- A. Add a second disk group with a new cache device and distribute the capacity disks equally.
- B. Add a second cache device to the same disk group to increase the caching size.
- C. Modify advanced vSAN configuration parameters to utilize the complete 800 GB for caching.
- D. Replace the existing cache device with a larger cache device.

Suggested Answer: C

Community vote distribution



jsi928 1 year, 1 month ago

Answer A. There is a 600GB Cap on Cache disks in vSAN 7.

upvoted 1 times

Lazylinux 2 years, 8 months ago

Selected Answer: A

is correct based on VMware recommended approach as D requires far more work,

It is relatively easy to scale up both cache and capacity together with the introduction of new disk groups. It is also easy to add additional capacity by inserting new agnetic disks to a disk group in hybrid (orflash devices for all-flash). But it could be much more difficult to add additional cache capacity. This is especially true if there is a need to swap out the current cache device and replace it with a newer larger one. Of course, this approach is also much more expensive. It is far easier to overcommit on flash resources to begin with rather than trying to increase it once vSAN is in production.

Design decision: Design with additional flash cache to allow easier scale up of the capacity layer. Alternatively scaling up cache and capacity at the same time through the addition of new disks groups is also an easier approach than trying to simply update the existing flash cache device in an existing disk group.

<https://core.vmware.com/resource/vmware-vsan-design-guide#sec6843-sub8>

upvoted 3 times

Lazylinux 2 years, 9 months ago

Selected Answer: D

Cache Must be 10% of capacity

upvoted 1 times

Lazylinux 2 years, 8 months ago

Ignore not quite correct - see my other comment

upvoted 1 times

bertrand14 3 years, 4 months ago

D is correct

upvoted 2 times

Rabbit117 3 years, 6 months ago

I think A.

upvoted 4 times

DenZn 3 years, 6 months ago

A should be correct

upvoted 2 times

An administrator receives an alert for vCenter being unavailable. With vCenter running on vSAN, the administrator wants to know if any other VMs are impacted.

What command can an administrator run to determine the overall health of the vSAN objects?

- A. esxcli vsan storage list
- B. esxcli vsan cmmnds timemachine get
- C. esxcli vsan cluster get
- D. esxcli vsan health cluster list

Suggested Answer: C

Reference:

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vmware-virtual-san-health-check-guide-6.1.pdf>
(66)

Community vote distribution

D (100%)

🗨️ **MosabSh** 2 years, 5 months ago

Selected Answer: D

D not doubt just run the command in ssh on esxi
upvoted 1 times

🗨️ **Lazylinux** 2 years, 9 months ago

Selected Answer: D

D not doubt just run the command in ssh on esxi
upvoted 2 times

🗨️ **wpjoulekkading** 3 years, 2 months ago

D definitely
upvoted 2 times

🗨️ **mparayil** 3 years, 4 months ago

D is correct agree
upvoted 2 times

🗨️ **Rabbit117** 3 years, 6 months ago

D is correct.
upvoted 2 times

🗨️ **DenZn** 3 years, 6 months ago

Agree - D should be correct
upvoted 2 times

🗨️ **Alefin** 3 years, 7 months ago

D is correct
upvoted 2 times

A failed storage controller has two vSAN disk groups attached. The components contained on the drives in those disk groups are marked Degraded. vSAN reports that some objects do not comply with their assigned storage policy. How is the compliance issue resolved?

- A. Initiate a proactive rebalance to force component repairs.
- B. Degraded components are repaired after CLOM Repair Delay Timer expires.
- C. Rebuilding of the degraded components starts immediately.
- D. Storage policies must be reapplied to all affected objects.

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Lazylinux** 2 years, 8 months ago

Selected Answer: C

is correct

upvoted 1 times

🗨️ 👤 **wpjoulekkading** 3 years, 2 months ago

C Definitely

upvoted 2 times

🗨️ 👤 **Rabbit117** 3 years, 5 months ago

C is correct.

upvoted 2 times

🗨️ 👤 **Alefin** 3 years, 6 months ago

C is correct

<https://cormachogan.com/2014/12/04/vsan-part-30-difference-between-absent-degraded-components/#:~:text=Since%20the%20disk%20may%20be,components%20are%20marked%20as%20DEGRADED.&text=If%20they%20do%20not%20become>

upvoted 2 times

🗨️ 👤 **DenZn** 3 years, 6 months ago

C is correct - <https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID-3863B640-3449-46A2-84E0-AC07E5A604FD.html>

upvoted 4 times

During planned maintenance of a four-node vSAN cluster, an outsourced IT contractor accidentally removed a 2.5" SSD cache disk from one of the vSAN nodes.

The storage policy has been configured with FTT=1 RAID 1, and the disk management UI marked the disk group as absent.

Which remediation steps should the administrator select to ensure VMs become compliant with the storage policy as soon as possible?

- A. replace the SSD cache disk > rescan > add disk back in disk group
- B. enter host in maintenance > remove from disk group > shutdown host > replace disk > power on host
- C. remove disk group > enter host in maintenance > fully evacuate all data > exit maintenance
- D. vSAN Health Check > retest > object health > repair objects immediately

Suggested Answer: A

🗨️ **Lazylinux** 2 years, 8 months ago

D is correct => FTT=1 RAID 1 only need 3 hosts but we have 4 hosts then we have 1 extra host to rebuild on
upvoted 1 times

🗨️ **MosabSh** 2 years, 9 months ago

D: a host will be absent for longer than 60 minutes. The affected VMs are still accessible however non-compliant with their storage policies. More importantly in the case of FTT=1, until a rebuild occurs vSAN will not be able to tolerate additional failures. If this is the case, you may choose to repair the objects immediately. This option will resynchronize the absent objects on available hosts in the vSAN cluster.
upvoted 1 times

🗨️ **wpjoulekkading** 3 years, 2 months ago

D no need to put host into maintenance mode or remove the disk group. Just put disk back into the host, do vSAN health check and click on the repair objects button.
upvoted 2 times

🗨️ **Lazylinux** 2 years, 8 months ago

Your explanation is wrong, as D dont mention put disk back!! now since we have 4 hosts and FTT=1 RAID 1 only need 3 hosts then we have 1 extra host to rebuild on and hence answer D
upvoted 1 times

🗨️ **Rabbit117** 3 years, 5 months ago

I think D is correct.

<https://blogs.vmware.com/virtualblocks/2017/11/09/understanding-vsan-rebuilds/#:~:text=Repair%20Objects%20Immediately&text=More%20importantly%20in%20the%20case,hosts%20in%20the%20vSAN%20cluster.>

upvoted 2 times

🗨️ **DSA** 3 years, 6 months ago

Should be D imho.

The question is to asap get the VMs compliant. Not on how the Disk is changed.

And the real question should be: Should this Administrator work in IT? ^^

upvoted 4 times

A 3-node All-Flash vSAN cluster has this configuration:

- ⇒ One disk group per host
- ⇒ All disk groups have identical make & model disks of an 800GB NVMe SSD cache and 2x 4TB SAS SSDs each

An administrator has been tasked with expanding storage capacity by adding one additional disk group with this configuration:

- ⇒ 1.6TB SAS SSD cache and 1x 4TB SAS SSDs

Assuming all the disks are compliant with the vSAN HCL, which is true?

- A. The additional disk group would not be possible as vSAN would not allow it.
- B. The additional disk group would work and be fully supported by VMware support.
- C. The additional disk group will generate an "Advanced vSAN Configuration in sync" health check warning.
- D. The additional disk group would make this cluster an unsupported configuration.

Suggested Answer: D

Reference:

<https://blogs.vmware.com/virtualblocks/2019/04/18/vsan-disk-groups/>

🗨️ **jsi928** 1 year, 1 month ago

You can't mix SSD and NVMe as there are vast performance differences. VMware does not support NVMe on Tri-Mode controllers:

<https://kb.vmware.com/s/article/88722> Answer is D.

upvoted 1 times

🗨️ **Lazylinux** 2 years, 8 months ago

B i would agree

upvoted 1 times

🗨️ **wpjoulekkading** 3 years, 2 months ago

I would pick B, it should work but definitely not best practise or recommended and i'm sure VMware Support would not support it either.

upvoted 1 times

🗨️ **Alefin** 3 years, 6 months ago

I think B

upvoted 1 times

🗨️ **dbus81** 3 years, 6 months ago

B it is supported but mostly not the best choice

upvoted 2 times

Which two statements are true about enabling deduplication and compression on an All-Flash vSAN cluster with data? (Choose two.)

- A. On-disk format version 3.0 or later is required to support deduplication and compression.
- B. This can be enabled through storage policies and does not require any data migration.
- C. Deduplication and Compression cannot be enabled on an existing vSAN datastore with data.
- D. Move one host at a time to maintenance mode and enable deduplication and compression.
- E. This requires data evacuation and disk format upgrade of all the disk groups.

Suggested Answer: AE

Reference:

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/GUID-BD30E413-F870-4C25-9507-123F59D0A4B8.html>

Community vote distribution

AE (100%)

🗳️ 👤 **Akesh7** 1 year, 6 months ago

Selected Answer: AE

AE is correct

upvoted 1 times

🗳️ 👤 **Lazylinux** 2 years, 8 months ago

Selected Answer: AE

I would agree

upvoted 1 times

🗳️ 👤 **wpjoulekkading** 3 years, 2 months ago

A & E correct

upvoted 2 times

🗳️ 👤 **dbus81** 3 years, 6 months ago

A and E

upvoted 2 times

🗳️ 👤 **Alefin** 3 years, 6 months ago

A-E are correct

It was tested

upvoted 2 times

🗳️ 👤 **ihjun** 3 years, 6 months ago

I think A and D

upvoted 1 times

🗳️ 👤 **ihjun** 3 years, 5 months ago

this is my mistake.. maintenance mode does not need. A&E correct

upvoted 2 times

🗳️ 👤 **Rabbit117** 3 years, 6 months ago

A and E are correct.

upvoted 2 times

Which two statements regarding networking are true for a vSAN Stretched Cluster? (Choose two.)

- A. A 5ms RTT or less is required between the data sites.
- B. The vSAN data network must communicate with the vSAN witness host.
- C. PIM is required for stretched vSAN clusters.
- D. A Stretched Cluster can use Layer 2 or Layer 3 communication for the vSAN network.
- E. A Stretched Layer 2 network is required at the vSAN data sites and at the witness site.

Suggested Answer: AD

Reference:

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

Community vote distribution

AD (100%)

🗳️ 👤 **Lazylinux** 2 years, 9 months ago

Selected Answer: AD

for sure

upvoted 1 times

🗳️ 👤 **wpjoulekkading** 3 years, 2 months ago

A & D correct

upvoted 2 times

🗳️ 👤 **dbus81** 3 years, 6 months ago

A and D

upvoted 2 times

🗳️ 👤 **Alefin** 3 years, 6 months ago

Definilty

upvoted 1 times

🗳️ 👤 **Rabbit117** 3 years, 6 months ago

A and D are correct.

upvoted 2 times

Which two storage policy changes require a full rebuild of all components for an object? (Choose two.)

- A. Increase the number of disk stripes for an object from 2 to 3.
- B. Increase the IOPS limit for an object.
- C. Reduce the Primary Failures to Tolerate from 2 to 1 for an object using a failure tolerance method of RAID-1 (Mirroring.)
- D. Disable object checksum on an object.
- E. Change the failure tolerance method of an object from RAID-1 (Mirroring) to RAID-5 (Erasure Coding).

Suggested Answer: CE

Reference:

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

 **Rabbit117** Highly Voted 3 years, 6 months ago

I think A and E are correct.

<https://cormachogan.com/2018/02/21/policy-changes-can-trigger-rebuild-vsan/?amp=1>

upvoted 7 times

 **wpjoulekkading** 3 years, 2 months ago

Agree i think A&E are correct

upvoted 3 times

 **jsi928** Most Recent 1 year, 1 month ago

A and E. Disabling checksum does not only Enabling does. Changing FTM does. <https://cormachogan.com/2018/02/21/policy-changes-can-trigger-rebuild-vsan/>

upvoted 1 times

 **Lazylinux** 2 years, 8 months ago

C and E for me ...Reduce the Primary Failures to Tolerate from 2 to 1 means from 5 hosts we now have 3 hosts formula $2n+1$ and hence total rebuild is required. As for E no need explain since all agree on....in Summary the FT method and tolerance level are two primary causes of SP rebuild

upvoted 1 times

 **Lazylinux** 2 years, 8 months ago

adding more info ..here vmware comments

Because the managed storage policy varies based on cluster size, adding or removing hosts will trigger a storage policy reconfiguration if it changes the size of the cluster so that it requires a different policy. For example, if you add an additional host to a cluster containing five hosts, the storage policy for that cluster is reconfigured from using 1 failure - RAID-1 (Mirroring) to 2 failures - RAID-6 (Erasure Coding). The reverse happens if the extra host is removed and the number of hosts is reduced from six to five.

<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 1 times

 **Alefin** 3 years, 6 months ago

C-E are correct

increase disk stripe needs resync data, not rebuild

upvoted 1 times

 **Alefin** 3 years, 6 months ago

A-E are correct

upvoted 2 times

An architect has been asked to plan the impact of a storage policy change to one of the production All-Flash based vSAN datastores. The change will reduce the failures to tolerate value from 2 to 1 for objects with a RAID-1 (Mirroring) layout in the existing policy. This change will be applied to multiple VMs.

Which statement is correct in this scenario?

- A. FTT2 requires another change to apply a RAID-6 configuration.
- B. The policy will need to be recreated and applied with an additional tag.
- C. There will be no objects rebuilt that impact performance.
- D. There will be an object rebuild impact on performance.

Suggested Answer: D

  **Rabbit117** Highly Voted 3 years, 6 months ago

I think C is correct. I have tested changing FTT from 1 to 0 and there was no component rebuild, I assume it's the same when going from 2 to 1.
upvoted 7 times

  **Lazylinux** Most Recent 2 years, 8 months ago

D for me remember going from FT2 to 1 meaning you losing 2 hosts from the cluster based on $2n+1$ and hence reconfig of the storage policy
..here vmware comments

Because the managed storage policy varies based on cluster size, adding or removing hosts will trigger a storage policy reconfiguration if it changes the size of the cluster so that it requires a different policy. For example, if you add an additional host to a cluster containing five hosts, the storage policy for that cluster is reconfigured from using 1 failure - RAID-1 (Mirroring) to 2 failures - RAID-6 (Erasure Coding). The reverse happens if the extra host is removed and the number of hosts is reduced from six to five.

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

upvoted 1 times

  **DSITTA** 2 years, 10 months ago

C is correct. D should be right if you have RAID5/6 that need a complete redistribution of coded parity; also need some additional space because of old copies are not removed until syncing new copies. With RAID 1, only witness and second copy are deleted. Non rebuild
upvoted 1 times

  **wpjoulekkading** 3 years, 2 months ago

C would be my pick. Changing a FTT from 2 to 1 would simply remove/delete additional disk components from disk to comply with the storage policy. there will be no rebuild/resync of data, only when you change the policy from FTT 1 to 2 there will be a performance impact.
upvoted 3 times

  **vJoeyB** 3 years, 3 months ago

C is correct
upvoted 2 times

  **dbus81** 3 years, 6 months ago

D is correct
upvoted 4 times

  **HappyPower** 3 years, 3 months ago

D is correct, I try it in my lab
upvoted 2 times

Which storage policy can be assigned to virtual machines in a 5-node, all flash vSAN cluster, to provide the highest level of redundancy?

- A. 1 failure "" RAID-5 (Erasure Coding)
- B. 2 failures "" RAID-6 (Erasure Coding)
- C. 2 failures "" RAID-1 (Mirroring)
- D. 3 failures "" RAID-1 (Mirroring)

Suggested Answer: B

Community vote distribution

C (100%)

 **Akesh7** 1 year, 6 months ago

Selected Answer: C

RAID1 with FTT2

upvoted 1 times

 **Lazylinux** 2 years, 8 months ago

Selected Answer: C

is correct

upvoted 1 times

 **MosabSh** 2 years, 9 months ago

Selected Answer: C

C is correct. 5 hosts support maximum 2FTT

upvoted 1 times

 **MosabSh** 2 years, 9 months ago

C is correct. 5 hosts support maximum 2FTT

upvoted 1 times

 **Rabbit117** 3 years, 5 months ago

C is correct. While A will enable redundancy to be maintained after a host failure, it does not protect against two simultaneous failures. C will give a higher level of redundancy.

upvoted 3 times

 **dbus81** 3 years, 6 months ago

C is correct , FTT 2n+1

upvoted 3 times

 **Rabbit117** 3 years, 6 months ago

I think A is correct. Using Raid 5 in a 5 node cluster will allow for redundancy to be maintained during maintenance or if a host fails.

upvoted 1 times

 **DenZn** 3 years, 6 months ago

C is correct - <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>

Raid 6 needs 6 Hosts

upvoted 3 times

An administrator observes this health check warning:

Controller firmware is VMware certified

They proceed to update the latest firmware on the Storage Controller. After a successful firmware upgrade, the same health check alert continues to display.

Which statement is true?

- A. Controller is not VMware certified and hence any firmware version displays a warning.
- B. The recently installed firmware is yet to be certified and included in VMware Compatibility Guide.
- C. The VMware vSAN Health Service needs to be restarted for changes to take effect.
- D. The alert can be ignored and silenced, since the firmware is newer than the recommended firmware.

Suggested Answer: *B*

Reference:

<https://kb.vmware.com/s/article/2150398>

  **Rabbit117** 3 years, 6 months ago

I think B is correct.

upvoted 2 times

Which is true of the vSAN iSCSI target service?

- A. It is a VMware supported cost effective way to provide shared storage to another hypervisor.
- B. The vSAN default storage policy determines the characteristics of the vSAN iSCSI target service storage objects.
- C. Security is provided through the iSCSI CHAP & Mutual CHAP Authentication methods.
- D. It is a VMware supported solution to provide iSCSI targets to applications that require multiple connections per session (MCS).

Suggested Answer: B

Community vote distribution

C (100%)

🗨️ 👤 **MosabSh** 2 years, 9 months ago

Selected Answer: C

Support for CHAP & Mutual CHAP authentication methods

https://kb.vmware.com/s/article/57344?lang=en_US

upvoted 1 times

🗨️ 👤 **dbus81** 3 years, 6 months ago

yes for C

upvoted 2 times

🗨️ 👤 **Rabbit117** 3 years, 6 months ago

C is correct.

upvoted 2 times

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

Answer: C https://kb.vmware.com/s/article/57344?lang=en_US

upvoted 3 times

🗨️ 👤 **DenZn** 3 years, 6 months ago

I think it is Answer C

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

upvoted 2 times

A vSAN administrator wants to implement end-to-end prioritization of vSAN traffic across the network in a shared network infrastructure that is using vSphere Distributed Switches (VDS). Which two can help achieve this objective? (Choose two.)

- A. Configure CoS or DSCP with high priority tag at the VDS and equivalent in the physical network.
- B. Enable jumbo frames for vSAN VMkernel ports and configure LACP for optimal load balancing.
- C. Enable Network I/O Control and allocate higher shares for vSAN traffic.
- D. Configure multiple vSAN VMkernel interfaces to load balance across multiple uplinks.
- E. Enable network resource pool at the VDS level to prioritize vSAN traffic.

Suggested Answer: BC

 **Parziwolle** Highly Voted 3 years, 5 months ago

I think A and C too:

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vsan-planning.doc/GUID-031F9637-EE29-4684-8644-7A93B9FD8D7B.html>

"Example Network I/O Control Configuration for a Physical Adapter That Handles vSAN"

"Marking vSAN Traffic"

upvoted 6 times

 **jsi928** Most Recent 1 year, 1 month ago

A and C. Jumbo frames do not prioritize traffic.

upvoted 1 times

 **Lazylinux** 2 years, 8 months ago

A and C, based on this article, same as one below

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vsan-planning.doc/GUID-031F9637-EE29-4684-8644-7A93B9FD8D7B.html>

upvoted 2 times

 **jebu1204** 3 years, 5 months ago

I think A and C: <https://blogs.vmware.com/virtualblocks/2019/04/21/designing-vsan-networks-2019-update/>

upvoted 4 times

 **Rabbit117** 3 years, 5 months ago

I think C and E are correct.

upvoted 1 times

 **Alefin** 3 years, 6 months ago

I think C-E

upvoted 1 times

A vSAN administrator is working on a 4 host homogenous vSAN Cluster that has 2 disk groups per host. One of the hosts needs to be permanently removed and repurposed for some other use.

These are the out of order steps which need to be completed:

1. Delete the disk group(s) that reside on the host you want to decommission.
2. Wait for resync traffic to complete and the host to enter maintenance mode.
3. Ensure there is sufficient capacity in the vSAN disk groups to decommission a node.
4. Put the host into maintenance mode with full data migration selected.
5. Move the ESXi host out of the cluster.

What is the correct order?

- A. 3, 1, 4, 2, 5
- B. 3, 4, 2, 1, 5
- C. 4, 3, 2, 5, 1
- D. 4, 3, 2, 1, 5

Suggested Answer: B

Reference:

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

Community vote distribution

B (100%)

🗨️ **Rabbit117** Highly Voted 3 years, 6 months ago

I agree with B.

upvoted 6 times

🗨️ **Akesh7** Most Recent 1 year, 6 months ago

Selected Answer: B

B is correct

upvoted 1 times

🗨️ **Jorge1235** 1 year, 6 months ago

Have I already taken the exam? I spend it with this dump?

upvoted 1 times

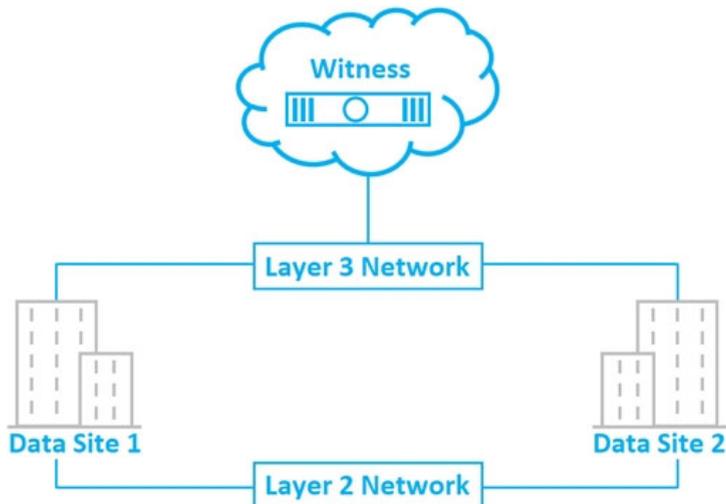
🗨️ **Lazylinux** 2 years, 8 months ago

Selected Answer: B

is correct

upvoted 1 times

Refer to the exhibit.



A vSAN Cluster stretched across Data Site 1 and Data Site 2 has a vSAN data network configured with an MTU of 9000. The witness node has been placed in a remote site across a WAN with an MTU of 1500. The vSAN health check, 'vSAN: MTU check (ping with large packet size)', reports a failed ping.

What two actions could resolve the issue? (Choose two.)

- A. Reduce the MTU from 9000 to 1500 on all vSAN VMkernel interfaces.
- B. Reconfigure the VLAN of the witness traffic to be identical to data nodes.
- C. Configure proper routing from data sites to the witness across the WAN.
- D. Reconfigure WAN link to support an MTU of 9000.
- E. Stretch L2 traffic from WAN to both data sites.

Suggested Answer: AD

Reference:

<https://kb.vmware.com/s/article/2141733>

Lazylinux 2 years, 8 months ago

A & D are correct

upvoted 1 times

Rabbit117 3 years, 6 months ago

I think A and D are correct.

<https://kb.vmware.com/s/article/2141733>

upvoted 1 times

Refer to the exhibit.

Overview

Object repair timer: 60 minutes ⓘ

This section is automatically refreshed every 10 seconds.

RESYNC THROTTLING

> Total resyncing objects	0
> Bytes left to resync	0.00 B
> Total resyncing ETA	0 seconds
> Scheduled resyncing	767 objects

An administrator is managing a 4-node, hybrid vSAN cluster. One of the hosts in the cluster went offline.

Referencing the information displayed in the resyncing objects section of the vSAN UI, what is the current state of this cluster?

- A. Offline host reconnected to the cluster and all objects are synchronized.
- B. Objects will be resynchronized when the Object Repair Timer expires.
- C. Resync Throttling has caused data resynchronization scheduling.
- D. Metadata resynchronization is scheduled to occur every 60 minutes.

Suggested Answer: D

Reference:

<https://core.vmware.com/resource/vsan-2-node-guide#sec11-sub3>

Community vote distribution

B (100%)

🗨️ **Rabbit117** Highly Voted 3 years, 6 months ago

I think B is correct.
upvoted 6 times

🗨️ **jsi928** Most Recent 1 year, 1 month ago

B is correct. <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vsan-monitoring.doc/GUID-57D74904-A8C4-4F3B-8D2A-1017AA8EBA85.html>
upvoted 1 times

🗨️ **Akesh7** 1 year, 6 months ago

Selected Answer: B

Correct is B
upvoted 2 times

🗨️ **Lazylinux** 2 years, 8 months ago

B is correct
upvoted 2 times

🗨️ **JoaoCoutinho** 3 years, 6 months ago

B: <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vsan-monitoring.doc/GUID-57D74904-A8C4-4F3B-8D2A-1017AA8EBA85.html>
upvoted 4 times

🗨️ **DenZn** 3 years, 6 months ago

I think Answer B should be correct
upvoted 3 times

An infrastructure architect is designing a vSAN Cluster that must satisfy these requirements:

- ⇒ support storage policies with 2 failures "" RAID-1 (Mirroring)
- ⇒ support storage policies with 2 failures "" RAID-6 (Erasure Coding) virtual machines must maintain compliance with the storage policy rules above when one host is in maintenance mode with full data migration

What is the minimum number of hosts needed to achieve these design requirements?

- A. 6
- B. 8
- C. 5
- D. 7

Suggested Answer: A

Reference:

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

🗨️ 👤 **Lazylinux** 2 years, 8 months ago

D for sure

upvoted 2 times

🗨️ 👤 **dbus81** 3 years, 6 months ago

Agree with D

upvoted 2 times

🗨️ 👤 **Rabbit117** 3 years, 6 months ago

FTT=2,RAID requires min 5 hosts.

FTT=2, Raid requires min 6 hosts.

For either policy to remain compliant after full evacuation of host an extra host is required.

Answer D is correct.

upvoted 3 times

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

The VMs must be compliant after one host failure/maintenance... so, it needs to have one extra host on the cluster. Answer D

upvoted 3 times

A higher level of resiliency has been requested for an existing vSAN Stretched Cluster. A key requirement will be the addition of a third site to protect against a potential double site failure of the two mirrored sites in the stretched cluster.
What solution would meet the resiliency requirement?

- A. Replicate to the third site at the VM level using vSphere Replication.
- B. Configure vSAN Stretched Cluster with FTT=2 to replicate to the third site.
- C. Add the third site to the existing vSAN Stretched Cluster.
- D. Replicate to the third site at the vSAN datastore level using Site Recovery Manager.

Suggested Answer: C

Community vote distribution

A (100%)

jsi928 1 year, 1 month ago

A

VMware Site Recovery Manager™ can be utilized with vSAN and vSphere Replication to orchestrate the recovery of multiple virtual machines.
upvoted 1 times

Lazylinux 2 years, 8 months ago

Selected Answer: A

is correct

upvoted 1 times

Lazylinux 2 years, 8 months ago

Replication from the stretched cluster to the disaster recovery site is facilitated by vSphere Replication. As mentioned previously, per-virtual machine RPOs for replication between two vSAN datastores can be as low as five minutes. Site Recovery Manager automates the failover and fail-back processes between the stretched cluster and the disaster recovery site.

upvoted 2 times

gmalan 3 years, 2 months ago

A is the righth answer.."replicate at VM Level"..vSphere Replication replicate only at the VM level
upvoted 2 times

gkoylu 3 years, 2 months ago

A is the right answer.

<https://core.vmware.com/resource/vsan-disaster-recovery#sec10254-sub1>

upvoted 2 times

vJoeyB 3 years, 3 months ago

The correct answer is A

upvoted 2 times

DenZn 3 years, 4 months ago

Agree with D - <https://core.vmware.com/resource/vsan-disaster-recovery#sec1-sub1>

upvoted 1 times

Rabbit117 3 years, 6 months ago

Apologies, I meant Site Recovery Manager (SRM).

upvoted 1 times

Rabbit117 3 years, 6 months ago

I think A is correct.

A stretched cluster has 2 data sites and 1 witness site. Replication would be required to replicate the VMS to another site, either vsan or other. Storage Replication manager orchestrates the failover of the replicated VMs, it requires replication be in place and working.

upvoted 3 times

An administrator is patching a production Stretched Cluster using VMware Update Manager.
Which statement is accurate?

- A. The witness hosts can be patched using normal patch processes with VMware Update Manager.
- B. Witness hosts cannot be patched or upgraded using VMware Update Manager.
- C. The witness host does not need to be upgraded as there is no data on the node, just metadata.
- D. The witness host must be deleted and a new witness host must be deployed.

Suggested Answer: A

🗨️ 👤 **Lazylinux** 2 years, 8 months ago

A is correct

Updating or Upgrading the vSAN Witness Appliance

The vSAN Witness Appliance can easily be maintained/patched using vSphere Update Manager in the same fashion as traditional ESXi hosts. It is not required to deploy a new vSAN Witness Appliance when updating or patching vSAN hosts. Normal upgrade mechanisms are supported on the vSAN Witness Appliance.

Note: When using an OEM provided vSphere ISO for upgrading vSAN hosts it is important to remember that additional OEM specific drivers or software may be included. It is important to use only a VMware provided vSphere ISO to upgrade the vSAN Witness Appliance.

<https://core.vmware.com/resource/vsan-stretched-cluster-guide#sec7306-sub2>

upvoted 1 times

🗨️ 👤 **dbus81** 3 years, 6 months ago

A is OK

upvoted 4 times

🗨️ 👤 **Rabbit117** 3 years, 6 months ago

I think A is correct, a witness host is a standard host.

upvoted 4 times

Which two are requirements for Kubernetes cluster virtual machines when leveraging vSphere Cloud Native Storage (CNS)? (Choose two.)

- A. configure UEFI firmware
- B. Set the Disk Mode to Independent
- C. set the disk.EnableUUID advanced setting to TRUE
- D. virtual machine hardware version 15 or higher
- E. minimum of 4 vCPUs per node

Suggested Answer: *CD*

Reference:

<https://docs.vmware.com/en/VMware-vSphere/6.7/Cloud-Native-Storage/GUID-BA795112-AFC4-4FCB-B5A6-1ACDCAB79ED3.html>

 **dbus81** 3 years, 6 months ago

C-D OK

upvoted 1 times

 **Alefin** 3 years, 6 months ago

C-D are correct

<https://docs.vmware.com/en/VMware-vSphere/6.7/Cloud-Native-Storage/GUID-BA795112-AFC4-4FCB-B5A6-1ACDCAB79ED3.html>

upvoted 2 times

 **Rabbit117** 3 years, 6 months ago

I agree with C and D.

upvoted 3 times

Which should be verified before enabling deduplication and compression?

- A. Adequate space is available.
- B. Ensure disk groups are balanced.
- C. Autoclaim mode is enabled.
- D. All hosts are in maintenance mode.

Suggested Answer: A

🗨️ 👤 **Lazylinux** 2 years, 8 months ago

A is correct as extra space is need for the formatting operation
upvoted 1 times

🗨️ 👤 **Rabbit117** 3 years, 6 months ago

I agree with A.
Additional space will be used to rebuild components as the disk groups are reformatted.
upvoted 4 times

A storage administrator is consolidating 2 data centers running on an HCI platform. There is a single site deployment of VMs running on the HCI platform with vSAN encryption. The administrator would like to protect data using replication and has these tools available:

- ⇒ vSphere replication
- ⇒ VM encryption
- ⇒ VMware vSAN

Considering the scenario, which statement is correct?

- A. Encrypted VM disks are not supported with VMware vSAN, only encrypted vSAN datastore.
- B. If a source VM has encrypted disks, the VM storage policy must be encrypted.
- C. If source VMs have encrypted disks, the target storage policy cannot include a VM encryption rule.
- D. If VMs have encrypted disks, the VM storage policy has a choice of encrypted or not encrypted.

Suggested Answer: A

Community vote distribution

B (100%)

🗨️ 👤 **Lazylinux** 2 years, 8 months ago

Selected Answer: B

I believe is correct

upvoted 1 times

🗨️ 👤 **Rabbit117** 3 years, 5 months ago

I agree with Alefin, I think B is correct.

When using VM Replication if the source VM is encrypted then you must select a storage policy at the target which has encryption enabled.

See link provided by Alefin.

upvoted 3 times

🗨️ 👤 **Alefin** 3 years, 6 months ago

I think B is correct

<https://docs.vmware.com/en/vSphere-Replication/8.3/com.vmware.vsphere.replication-admin.doc/GUID-25086430-782F-41F9-A7C1-4AF96668597E.html>

upvoted 1 times

Three months ago, an administrator deployed vRealize Operations Manager from vCenter to manage his newly deployed four-node vSAN 6.7 All-Flash cluster.

What must the administrator do to continue using vRealize Operations Dashboards for vSAN?

- A. Reconfigure the vSAN cluster.
- B. Upgrade vRealize Operation Manager.
- C. Apply the correct license to the vSAN cluster.
- D. Reconfigure the vRealize Operations Manager.

Suggested Answer: C

jsi928 1 year, 1 month ago

Answer is D.

Set the number of days that vRealize Operations keeps the data, after which it is purged from the system. The default value is 90 days.

<https://docs.vmware.com/en/vRealize-Operations/8.10/com.vmware.vcom.user.doc/GUID-123DDD0F-DB86-4D4B-8AC1-A55C7D54276F.html>

upvoted 1 times

Lazylinux 2 years, 8 months ago

C is correct

upvoted 1 times

gmalan 3 years, 2 months ago

C es correct answer. "3 months ago" (90 day trial period has expired)

upvoted 3 times

DenZn 3 years, 4 months ago

Should be D

<https://docs.vmware.com/en/vRealize-Operations-Manager/7.5/com.vmware.vcom.core.doc/GUID-D2031BAE-43EA-47AA-AF53-1D62C86C3EA5.html>

upvoted 2 times

Rabbit117 3 years, 6 months ago

I think C is correct.

upvoted 2 times