



**Free Questions for NS0-520 by certscare**

**Shared by Flowers on 12-12-2023**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

## Question Type: MultipleChoice

---

An administrator has migrated an ISCSI LUN with Foreign LUN Import (FLI) tool. The migrated LUN is presented on the new system, but it is not visible to the host. In this scenario which additional action

must be completed to make the LUN visible?

### Options:

---

- A- Bind the igroup to a port set
- B- Install ANA software on the host
- C- Stop the FCP service
- D- Map the LUN to an ISCSI igroup

### Answer:

---

D

### Explanation:

---

Installing ANA is related with NVMe protocol

Binding an igroup to a port set will only limit our paths

Stopping the FCP service will not affect at anything, because this is an ISCSI LUN

When you migrate ISCSI using FLI, you need to map the LUNs to an igroup where the destination arrays is added, after the LUN has been imported, you can create a new iSCSI igroup and add the hosts to

the igroup

You can find more information on this link:

<https://www.netapp.com/pdf.html?item=/media/17060-tr4380pdf.pdf> (page 71)

## Question 2

---

**Question Type: MultipleChoice**

---

You are configuring a VMware sphere host for FC access to a Netapp ONTAP cluster. In this scenario, which two tasks does Netapp recommend performing on the Vsphere host? (choose two)

## Options:

---

- A- Set the queue depth on the host to 8
- B- Update the HBA driver
- C- Zone the host by using the host WWNNs
- D- Use the Netapp Virtual Storage Console to configure the host settings

## Answer:

---

B, D

## Explanation:

---

Zoning the host using WWNNs is not recommended. Netapp always recommends to use WWPN to zone correctly.

We can not establish a queue depth value because we don't have the amount of ports and hosts to calculate the recommended queue depth

Installing VSC and updating the HBA driver are best practices recommended by Netapp

You can find more information on this link:

<https://docs.netapp.com/ontap-9/index.jsp?topic=%2Fcom.netapp.doc.exp-iscsi-esx-cpg%2FGUID-C751F905-04B4-4DE6-ADC4-A47EA9076078.html>

<https://docs.netapp.com/ontap-9/topic/com.netapp.doc.exp-fc-esx-cpg/FC%20express%20configuration%20for%20ESXi%20using%20VSC.pdf> (page 7)

## Question 3

---

**Question Type:** MultipleChoice

---

You want to use ANA on SUSE Enterprise Linux. Which two components need to be verified in this scenario? (choose two)

### Options:

---

- A- Verify the version of ONTAP
- B- Verify the version of SUSE Enterprise Linux
- C- Verify that ALUA must be enabled
- D- Verify that the ANA driver is installed

### Answer:

---

A, B

## Explanation:

---

ANA Supportability

NVMe/FC is supported on ONTAP 9.6 or later for the following versions of SLES:

SLES15 SP1

SLES15 SP1 host can run both NVMe/FC, & FCP traffic through the same fibre channel initiator adapter ports.

You can find more information on this link:

[https://docs.netapp.com/us-en/ontap-sanhost/nvme\\_sles15\\_sp1.html#supportability](https://docs.netapp.com/us-en/ontap-sanhost/nvme_sles15_sp1.html#supportability)

## Question 4

---

**Question Type:** MultipleChoice

---

You have a Netapp ONTAP cluster and are configuring ISCSI for some new hosts . You want to ensure that ISCSI connectivity can survive physical network failures.

In this scenario, which two features would help fulfill this requirement? (choose two)

**Options:**

---

- A- Broadcast domains
- B- Failover groups
- C- MPIO
- D- Interface groups

**Answer:**

---

C, D

**Explanation:**

---

Broadcast Domains are used to segment and create failover groups automatically... Broadcast Domain do not guarantee redundancy on ISCSI because failover groups are not used on ISCSI.

- Failover groups do not apply in SAN **iSCSI** or FC environments.
- MCS is not supported with clustered Data ONTAP. Since MPIO can be used for clustered Data ONTAP or Data ONTAP running in 7-Mode, NetApp recommends using MPIO consistently across architectures.
- MCS is an iSCSI-only feature. Since MPIO can be used for FC and/or iSCSI, NetApp recommends using MPIO consistently across protocols.
- MCS is a feature for Windows only. Since MPIO can be used for all host operating systems supported by NetApp, NetApp recommends using MPIO consistently across host operating systems.

## Question 5

---

### Question Type: MultipleChoice

---

You created a LUN in a new volume on your Netapp ONTAP cluster. You enabled the default snapshot schedule and enabled snapshot autodelete on the volume. In this scenario, in which three situations

does the snapshot autodelete feature automatically delete snapshot copies in the volume? (choose three)



### Options:

---

- A- The snapshot reserve space is nearly full
- B- The LUN is nearly full
- C- The overwrite reserve space is full
- D- The volume is nearly full
- E- The aggregate is nearly full

### Answer:

---

A, C, D

### Explanation:

---

Depending on what is selected in the 'Trigger' option, Snapshot Autodelete will automatically delete a snapshot when, either the volume space utilized, or the snapshot reserve reaches a threshold capacity

specified below.

Automatic deletion of Snapshot copies from flexible volumes Data ONTAP can automatically delete one or more Snapshot copies on the volume, provided the Data ONTAP Snapshot copy autodeletion

policy is enabled and set to trigger when the overwrite reserve is nearly full on the volume.

You can find more information on this link:

[https://kb.netapp.com/Advice\\_and\\_Troubleshooting/Data\\_Storage\\_Software/ONTAP\\_OS/How\\_to\\_use\\_Snapshot\\_Autodelete](https://kb.netapp.com/Advice_and_Troubleshooting/Data_Storage_Software/ONTAP_OS/How_to_use_Snapshot_Autodelete)

<https://library.netapp.com/ecmdocs/ECMP1217281/html/GUID-98FB8EC4-F3AE-4D25-AF44-E3200AFEB1AB.html>

## Question 6

---

**Question Type: MultipleChoice**

---

File deletions from a SAN-attached host that is using Windows Server 2016 are not freeing up space in LUN lun3. Referring to the exhibit, which setting is preventing the SCSI UNMAP operation?

### Options:

---

- A- Space reservation honored is set to false
- B- Space reservation is set to disabled
- C- Clone Autodelete Enabled is set to false
- D- Space Allocation is set to disabled

**Answer:**

---

D

**Explanation:**

---

For clustered Data ONTAP, support for SCSI UNMAP begins in the 8.2 release family. For a LUN to advertise support for and accept SCSI UNMAP commands, enable the space-allocation option on LUNs which you expect to use this feature with.

You can find more information on this link:

[https://kb.netapp.com/Advice\\_and\\_Troubleshooting/Data\\_Storage\\_Software/ONTAP\\_OS/ONTAP\\_and\\_Windows\\_SCSI\\_unmap\\_and\\_space\\_reclam](https://kb.netapp.com/Advice_and_Troubleshooting/Data_Storage_Software/ONTAP_OS/ONTAP_and_Windows_SCSI_unmap_and_space_reclam)

cluster1::> lun show -instance

Vserver Name: vs1  
LUN Path: /vol/host3/lun3  
Volume Name: host3  
Qtree Name: ""  
LUN Name: lun3  
LUN Size: 10.00GB  
OS Type: windows\_2008  
Space Reservation: disabled  
Serial Number: 21AFF?P4Fx08  
Serial Number (Hex): 5a6c4146463f503446783038  
Comment:  
Space Reservations Honored: false  
Space Allocation: disabled  
State: online  
LUN UUID: d23a1356-68e8-4770-b4c6-0bc374d5d5b9  
Mapped: mapped  
Physical Size of Logical Block: 512B  
Device Legacy ID: -  
Device Binary ID: -  
Device Text ID: -  
Read Only: false  
Fenced Due to Restore: false  
Used Size: 0  
Maximum Resize Size: 502.0GB  
Creation Time: 3/17/2020 18:37:12  
Class: regular  
Node Hosting the LUN: cluster1-01  
QoS Policy Group: -  
QoS Adaptive Policy Group: -  
Caching Policy Name: -  
Clone: false  
Clone Autodelete Enabled: false  
Inconsistent Import: false  
Application: -

## Question 7

---

**Question Type:** MultipleChoice

---

A storage administrator needs to ensure that an ONTAP FC LUN is accessible for a remote data center. The ISL ports are reporting errors. The trunkshow command output from a Brocade switch is shown below.

```
switch01:admin> trunkshow 1: 10-> 10 10:00:50:eb:1a:20:83:96 5 deskew 15 MASTER 2: 11-> 11 10:00:50:eb:1a:20:83:96  
deskew 16 MASTER
```

### Options:

---

- A- Set the ISL ports as FL\_Ports
- B- Set the ISL ports as N\_Ports
- C- Set the ISL ports as F\_Ports
- D- Set the ISL ports as E\_Ports

### Answer:

---

D

## **Explanation:**

---

On each switch fabric, you must configure the switch ports that connect the Inter-Switch Link (ISL). These ISL ports are otherwise known as the E-ports.

You can find more information on this link:

<https://library.netapp.com/ecmdocs/ECMP1636017/html/GUID-B2BBB025-02A9-442C-9910-00FEF253CE52.html>

## **Question 8**

---

### **Question Type: MultipleChoice**

---

During testing, an administrator notices that the wrong subsystem was returned using the NVMe discover command. The administrator was expecting test0\_subsystem as the subsystem

Referring to the exhibit, what would cause the discrepancy?

```

>vserv̄er nvme subsystem show
Vserver Subsystem      Target NQN
-----
test0    test0_subsystem nqn.1992-08.com.netapp:sn.87b3c578feeb11e794f200a098b3f653:subsystem.test0_subsystem
test1    test1_subsystem nqn.1992-08.com.netapp:sn.5b584a86fa6411e794f200a098b3f653:subsystem.test1_subsystem
test2    test2_subsystem nqn.1992-08.com.netapp:sn.12fea223fa6411e794f200a098b3f653:subsystem.test2_subsystem
test3    test3_subsystem nqn.1992-08.com.netapp:sn.288aa4f6feeb11e794f200a098b3f653:subsystem.test3_subsystem

# nvme discover --transport=fc -traddr=nn-0x203f00a098b3f7a7:pn0x204000a098b3f7a7 --host-traddr=nn-0x200000109b1c0x100000109b1c0f8f
Discovery Log Number of Records 1, Generation counter 17
====Discovery Log Entry 0====
trtype: fibre-channel
adrfam: fibre-channel
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: none
subnqn: nqn.1992-08.com.netapp:sn.5b584a86fa6411e794f200a098b3f653:subsystem.test1_subsystem
traddr: nn-0x203f00a098b3f7a7:pn-0x204000a098b3f7a7

```

## Options:

---

- A- The wrong NQN was mapped to the subsystem
- B- The wrong WWPN was used in the discovery command
- C- Asymmetric Namespace Access (ANA) was not properly configured
- D- The NVMe was not configured on the test0 SVM

## Answer:

---

A

## Explanation:

---

We can discard that ANA was not properly configured because this is not a problem about multipathing.

We can discard WWPN because this is an NVMe configuration where WWPN is not involved

We can discard that NVMe was not configured because we wouldn't be able to even run that command if that's true

Configure the NQN correctly is one of the steps when we configure NVMe

You can find more information on this link:

<http://senthil-it-infrastructure.blogspot.com/2020/05/netapp-ontap-97-nvme-configuration-and.html>

## Question 9

---

### Question Type: MultipleChoice

---

After performing failover testing on a newly deployed MetroCluster configuration, some aggregates remain in a degraded state. Referring to the exhibit, which statement is true?



```

cluster_A::> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes          RAID Status
-----
node_A_1data01_mirrored
    4.15TB      3.40TB    18% online    3 node_A_1      raid_dp,mirrored,normal
node_A_1root
    707.7GB    34.29GB    95% online    1 node_A_1      raid_dp,mirrored,normal
node_A_2_data01_mirrored
    4.15TB      4.12TB     1% online    2 node_A_2      raid_dp,mirror,degraded
node_A_2_data02_unmirrored
    2.18TB      2.18TB     0% online    1 node_A_2      raid_dp,normal
node_A_2_root
    707.7GB    34.27GB    95% online    1 node_A_2      raid_dp,mirror,degraded

```

```

cluster_A::> storage aggregate show-status -node node_A_2 !+root
Owner Node: node_A_2
Aggregate: node_A_2_data01_mirrored (online, raid_dp, mirror degraded) (block checksums)
  Plex: /node_A_2_data01_mirrored/plex0 (online, normal, active, pool0)
  RAID Group /node_A_2_data01_mirrored/plex0/rg0 (normal, block checksums)
...

```

```

Plex: /node_A_2_data01_mirrored/plex4 (offline, failed, inactive, pool1)
  RAID Group /node_A_2_data01_mirrored/plex4/rg0 (partial, none checksums)

```

Position	Disk	Pool	Type	RPM	Usable Size	Physical Size	Status
dparity	FAILED	-	-	-	827.7GB	-	(failed)
parity	FAILED	-	-	-	827.7GB	-	(failed)
data	FAILED	-	-	-	827.7GB	-	(failed)
data	FAILED	-	-	-	827.7GB	-	(failed)
data	FAILED	-	-	-	827.7GB	-	(failed)

```

Aggregate: node_A_2_data02_unmirrored (online, raid_dp) (block checksums)
  Plex: /node_A_2_data02_unmirrored/plex0 (online, normal, active, pool0)
  RAID Group /node_A_2_data02_unmirrored/plex0/rg0 (normal, block checksums)
...

```

15 entries were displayed.

## Options:

---

- A- The MetroCluster Tiebreaker software is misconfigured
- B- The ONTAP Mediator service is misconfigured
- C- A cluster switch is powered off
- D- A disk shelf is powered off

**Answer:**

---

D

**Explanation:**

---

The aggregate with drives on the powered-off shelf should have a degraded RAID status, and drives on the affected plex should have a failed status

You can find more information on this link:

[https://docs.netapp.com/us-en/ontap-metrocluster/install-ip/task\\_test\\_the\\_mcc\\_configuration.html#verifying-operation-after-loss-of-a-single-storage-shelf](https://docs.netapp.com/us-en/ontap-metrocluster/install-ip/task_test_the_mcc_configuration.html#verifying-operation-after-loss-of-a-single-storage-shelf)

## Question 10

---

**Question Type:** MultipleChoice

---

Your customer asks you how to connect a host for SAN access to the newly installed AFF A400 system. In this scenario, which two statements are correct? (choose two)

### Options:

---

- A- Direct Connection to the nodes from the hosts with ISCSI is supported
- B- Direct Connection to the nodes from the hosts with ISCSI is not supported
- C- Direct Connection to the nodes from the hosts with FC is not supported
- D- Direct Connection to the nodes from the hosts with FC is supported

### Answer:

---

A, C

### Explanation:

---

ISCSI-Direct-attachment: In a direct-attached configuration, one or more hosts are directly connected to the controllers.

You cannot directly attach FC or FC-NMVE SAN hosts to single nodes without using an FC switch.

You can find more information on this link:

## Question 11

---

### Question Type: MultipleChoice

---

While changing the network connections on your ONTAP cluster from twinax to fiber, the ports experience network connectivity issues. You want to verify which speeds the ports support and ensure that you have supported transceivers. In this scenario which two actions would accomplish this task (choose two)

### Options:

---

- A- Use hardware universe to determine the supported transceivers
- B- Use the interoperability matrix tool to determine the supported port speeds
- C- Use hardware universe to determine the supported port speeds
- D- Use the interoperability matrix tool to determine the supported transceivers

**Answer:**

---

A, C

**Explanation:**

---

Hardware universe will show information about port speeds and transceiver compatibility

You can find more information on this link:

[https://hwu.netapp.com/Resources/hwu\\_ug.pdf](https://hwu.netapp.com/Resources/hwu_ug.pdf) (page 73)

**To Get Premium Files for NS0-520 Visit**

**<https://www.p2pexams.com/products/ns0-520>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/netapp/pdf/ns0-520>**

