



Free Questions for NS0-520 by dumpshq

Shared by Fowler on 20-10-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You want to move a SAN LIF and preserve its configuration information and avoid rezoning the switch fabric. In this scenario, which two steps must be performed to fulfill this goal? (Choose two)

Options:

- A- Move the SAN LIF to another node in a cluster on a different storage virtual machine (SVMs)
- B- Keep the SAN LIF online before it is moved
- C- Take the SAN LIF offline before it is moved
- D- Move the SAN LIF to another node in a cluster on existing storage virtual machine (SVMs)

Answer:

C, D

Explanation:

Moving SAN LIFs

If a node needs to be taken offline, you can move a SAN LIF to preserve its configuration information, such as its WWPN, and avoid rezoning the switch fabric. Because a SAN LIF must be taken offline

before it is moved, host traffic must rely on host multipathing software to provide nondisruptive access to the LUN. You can move SAN LIFs to any node in a cluster, but you cannot move the SAN LIFs

between storage virtual machines (SVMs).

You can find more information on this link:

<https://docs.netapp.com/ontap-9/index.jsp?topic=%2Fcom.netapp.doc.dot-cm-sanag%2FGUID-09F5C2A3-FA66-4252-BEEF-D556A143CD30.html>

Question 2

Question Type: MultipleChoice

You upgraded your Linux hosts to SUSE Enterprise Linux 15. In this scenario, which two actions should you perform before you test asymmetric namespace access (ANA)? (Choose two)

Options:

- A-** Verify if the NVMe service has started
- B-** Verify if the LUNs and namespaces are in the same volume.
- C-** Verify if the version of Netapp ONTAP software is 9.5 or later
- D-** Verify if there is only one LIF in the SVM


Answer:

A, C

Explanation:

ONTAP 9.5 Introduces ANA

Figure 12) NVMe/FC Storage Failover: ONTAP 9.5 introduces ANA.



<https://www.netapp.com/de/media/10681-tr4684.pdf>

Question 3

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:

Question 4

Question Type: MultipleChoice

You are testing FCP paths failures on a 2-node Netapp AFF All SAN Array and verify that persistent ports are correctly enabled. Before testing, a host sees four optimized paths to a LUN. You perform a

node takeover and re-check the host.

Which statement is true after the node takeover?

Options:

- A- The host sees four optimized paths to the LUN
- B- The host sees two non-optimized paths and two dead paths to the LUN
- C- The host sees two optimized paths and two non-optimized paths to the LUN
- D- The host sees two optimized paths and two dead paths to the LUN

Answer:

A

Explanation:

Support for persistent ports

Beginning in ONTAP 9.8, persistent ports are enabled by default on All SAN Arrays (ASAs) that are configured to use the FC protocol. Persistent ports are only available for FC and require zone

membership identified by World Wide Port Name (WWPN).

Persistent ports reduce the impact of takeovers by creating a shadow LIF on the corresponding physical port of the HA partner. When a node is taken over, the shadow LIF on the partner node assumes the

identity of the original LIF, including the WWPN. Before the status of path to the taken over node is changed to faulty, the shadow LIF appears as an Active/Optimized path to the host MPIO stack, and I/O is

shifted. This reduces I/O disruption because the host always sees the same number of paths to the target, even during storage failover operations.

You can find more information on this link:

<https://docs.netapp.com/allsan/index.jsp?topic=%2Fcom.netapp.doc.dot-asa-config%2FGUID-646B3CFD-9E00-491A-A02E-F1668C5C9DBA.html>

Question 5

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 (page 73)

Question 6

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)

      Usable Physical
      Size      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

Question 7

Question Type: MultipleChoice

You have an existing 8-node Netapp Cluster serving NFS data, and you want to expand your cluster by an additional four nodes to serve ISCSI. In this scenario, where would you validate this configuration?

Options:

- A- Active IQ Upgrade Advisor
- B- Interoperability Matrix Tool
- C- SAN Administration Guide
- D- Hardware Universe

Answer:

D

Explanation:

Hardware Universe shows how many nodes and which filers are compatible to mix between them

Question 8

Question Type: MultipleChoice

In an NVMe/FC environment, the host cannot discover the namespace on the AFF system. Referring to the exhibit, what is the problem?

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd

*> vserver nvme subsystem host show -vserver vs_nvme_10
Vserver Subsystem Host NQN
-----
rhel_141_nvme_ss_10_0
nqn.2014-08.org.nvmexpress:uuid:3ca559e1-5588-4fc4-b7d6-5ccfb0b9f054
```

Options:

- A- The host MPIO drivers are not installed
- B- NVMe is not supported on the FC adapter
- C- The NQN is not correct
- D- The NVMe service is not running on the SVM

Answer:

C

Explanation:

We would not be able to run that command if NVMe is not enabled

NVMe is supported on the same SVM along with FC protocol

The multipath tool on NVMe is called ANA

NQN needs to be correctly declared on the subsystem to map the new NVMe resources

Question 9

Question Type: MultipleChoice

You install an ASA AFF A400 system with 32 GB FC SFP+ modules in port 0e and 0f. After creating the SVM and configuring the FCP protocol, the physical port stays offline. Referring to the exhibit which statement is correct?

```
cluster1::> system hardware unified-connect show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
cluster1-01	0e	cna	target	-	-	offline
cluster1-01	0f	cna	target	-	-	offline

Options:

- A- The current type of the ports must be configured as an initiator
- B- The current mode to the ports must be configured as an FC port
- C- The current mode for the ports must be set as an NVMe/FC port
- D- The admin status is manually set to down

Answer:

B

Explanation:

Before using your UTA ports on a FC infrastructure, you need to modify the current mode from CNA to FC. You need to set the ports offline before you modify that parameter.

You can find more information on this link:

<http://docs.netapp.com/ontap-9/index.jsp?topic=%2Fcom.netapp.doc.dot-cm-sanag%2FGUID-E2CD31BE-28EC-468E-AF49-83A641575799.html>

Question 10

Question Type: MultipleChoice

You are designing a new 4-node AFF A400 SAN cluster with 20 ISCSI hosts. Each AFF A400 node will have four data LIFs. Each ISCSI host needs to access every LUN over two LIFs per node on every

SAN node.

In this scenario, which Netapp SAN feature enable this connectivity?

Options:

A- Port sets

B- HBA queue depth

C- Selective LUN mapping

D- Asymmetric namespace

Answer:

A

Explanation:

SLM will choose 4 LIFs per node, being 8 LIF total (owner and HA partner), enabling portsets will help to limit the ports per node... that way we can establish only 2 LIF per node.

Creating port sets and binding igroups to port sets

In addition to using Selective LUN Map (SLM), you can create a port set and bind the port set to an igroup to further limit which LIFs can be used by an initiator to access a LUN. If you do not bind a port set

to an igroup, then all of the initiators in the igroup can access mapped LUNs through all of the LIFs on the node owning the LUN and the owning node's HA partner.

You can find more information on this link:

<https://docs.netapp.com/ontap-9/index.jsp?topic=%2Fcom.netapp.doc.dot-cm-sanag%2FGUID-5CC27202-A43F-429F-AB13-C70E7CD58E09.html>

Question 11

Question Type: MultipleChoice

Referring to the exhibit, which two actions would you perform to ensure that each host has valid multipathing? (choose two)

```
node      port  link broadcast-domain
-----
Cluster1-01 ala-10 up Storage
Cluster1-02 a2a-10 up Storage
```

```
Vserver      Logical      Status      Network      Current      Current Is
Interface    Admin/Oper  Address/Mask  Node          Port         Home
-----
Storage_iSCSI
Storage_iSCSI_lif1 up/up 10.10.10.30/24 Cluster1-01  ala-10  true
Storage_iSCSI_lif2 up/up 10.10.10.31/24 Cluster1-02  a2a-10  true
```

```
Vserver      Igroup  Protocol  OS Type  Initiators
-----
Storage_iSCSI vmware  iscsi    vmware  iqn.1998-01.com.vmware:5dfa0d08-9a33-fbc6-abac-94c691140409-59a59ccf
iqn.1998-01.com.vmware:5dfa0efa-8c55-5a44-a125-94c6911407fd-5972503d
```

```
Vserver      Tpgroup Name  TSIH Initiator Name
-----
Storage_iSCSI Storage_iSCSI_lif1 1 iqn.1998-01.com.vmware:5dfa0efa-8c55-5a44-a125-94c6911407fd-5972503d
Storage_iSCSI Storage_iSCSI_lif1 2 iqn.1998-01.com.vmware:5dfa0efa-8c55-5a44-a125-94c6911407fd-5972503d
Storage_iSCSI Storage_iSCSI_lif2 1 iqn.1998-01.com.vmware:5dfa0efa-8c55-5a44-a125-94c6911407fd-5972503d
Storage_iSCSI Storage_iSCSI_lif2 2 iqn.1998-01.com.vmware:5dfa0efa-8c55-5a44-a125-94c6911407fd-5972503d
Storage_iSCSI Storage_iSCSI_lif2 3 iqn.1998-01.com.vmware:5dfa0d08-9a33-fbc6-abac-94c691140409-59a59ccf
Storage_iSCSI Storage_iSCSI_lif2 4 iqn.1998-01.com.vmware:5dfa0d08-9a33-fbc6-abac-94c691140409-59a59ccf
6 entries were displayed.
```

Options:

- A- Verify that the VLANs are properly configured on the switches
- B- Verify that the VLANs are properly configured on the ONTAP cluster
- C- Verify that ALUA is enabled on the switches
- D- Verify that ISCSI service is started on the ONTAP cluster

Answer:

A, B

Explanation:

Iscsi is already started, we can not run those commands successfully if the protocol is not started

ALUA is enabled on hosts, It is not enabled on switches

VLANs are configured on the Netapp and switches, they need to be properly configured on both sides to work correctly

You can find more information on this link:

<https://docs.netapp.com/ontap-9/index.jsp?topic=%2Fcom.netapp.doc.dot-cm-nmg%2FGUID-52AE6E9C-13B6-49BA-B030-545C1C7647AC.html>

Question 12

Question Type: MultipleChoice

What are two benefits of implementing VLANs for ISCSI configuration? (choose two)

Options:

- A- Reduced number of broadcast domains
- B- Increased performance
- C- Increased security
- D- Increased resiliency

Answer:

A, C

Explanation:

When you implement VLANs in large IP network infrastructures, you derive the following benefits:

Increased security.

VLANs enable you to leverage existing infrastructure while still providing enhanced security because they limit access between different nodes of an Ethernet network or an IP SAN.

You can find more information on this link:

<https://docs.netapp.com/ontap-9/index.jsp?topic=%2Fcom.netapp.doc.dot-cm-sanconf%2FGUID-C5288E55-DAED-4050-84A2-71BF13BC6556.html>

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>

