

DUMPSBOSS.

NetApp Implementation Engineer - SAN ONTAP Exam

Netapp NS0-522

Version Demo

Total Demo Questions: 12

Total Premium Questions: 122

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, SAN Solution Design	40
Topic 2, SAN Solution Implementation	8
Topic 3, SAN Solution Configuration	53
Topic 4, SAN Solution Testing and Troubleshooting	21
Total	122

QUESTION NO: 1

What configuration must be applied for NVMe/FC?

- A. Create multiple initiator zones and multiple target zones.
- B. Configure igroup and IQN mapping.
- C. Enable NPV on all fabric switches.
- D. Enable NPIV on all fabric switches.

ANSWER: D

Explanation:

Enable NPIV on all fabric switches. NVMe/FC on ONTAP uses Fibre Channel connectivity, and N_Port ID Virtualization (NPIV) is a required fabric capability for the supported NVMe/FC design. NPIV permits a physical FC port to present multiple virtual N_Port identities (virtual WWPNNs). This is essential because ONTAP FC/NVMe target logical interfaces use virtual identities that must be able to log in to, and be recognized by, the Fibre Channel fabric independently of the physical adapter port. Enabling NPIV across the fabric path ensures that these virtual ports can be registered, zoned, and routed correctly between NVMe hosts and the ONTAP HA pair. The configuration should be applied consistently on every applicable fabric switch so that virtual-port logins remain supported through the end-to-end SAN fabric, including during path changes and failover events. NetApp's SAN configuration documentation identifies the supported host, protocol, and fabric requirements for ONTAP SAN deployments, while the ONTAP SAN administration documentation covers NVMe namespace and subsystem configuration: [NetApp ONTAP SAN configuration documentation](#) and [NetApp ONTAP SAN administration documentation](#).

QUESTION NO: 2

A four-path FCP configuration is deployed on a standard NetApp AFF system rather than an All SAN Array. The LUN is owned by node 1, and the host has two paths through node 1 and two paths through its HA partner, node 2.

Which ALUA path state should the host normally report?

- A. 4 Active/Optimized, 0 Active/Unoptimized
- B. 2 Active/Optimized, 2 Active/Unoptimized
- C. 0 Active/Optimized, 4 Active/Unoptimized
- D. 2 Active/Optimized, 0 Active/Unoptimized

ANSWER: B

Explanation:

On a standard AFF system, paths through the node that owns the LUN are normally reported as Active/Optimized. Paths through the HA partner remain available for nondisruptive access but are reported as Active/Unoptimized. Therefore, the host should report two Active/Optimized paths and two Active/Unoptimized paths.

An ASA system provides symmetric active-active access and would report all supported paths as optimized. That behavior does not apply to the standard AFF configuration in this scenario.

QUESTION NO: 3

A NetApp AFF node in a four-node cluster has a single FC port on one node, which is used for legacy tape backup. Tape is no longer needed, and the port is disconnected and modified for reuse in the FCP SAN.

Based on this change, which command will verify that the port is usable?

- A. network interface show
- B. lun persistent-reservation show
- C. system node hardware unified-connect show
- D. vserver fcp initiator show

ANSWER: C

Explanation:

`system node hardware unified-connect show` is the appropriate command to verify the operational configuration of a Unified Target Adapter or unified-connect port after it has been repurposed for SAN use. ONTAP uses unified-connect hardware settings to determine whether supported FC/CNA-capable adapters are configured for their intended protocol and personality. The command displays adapter-level information, including the current and pending configuration, allowing the administrator to confirm that the port is configured for Fibre Channel operation and is available for use by the FCP SAN service. This validation is especially relevant after a physical disconnect or adapter reconfiguration, because a port can require a configuration change and, in some cases, a reboot or other activation step before its new mode becomes active. Confirming the current state from the node ensures the hardware is ready before proceeding with SAN configuration tasks such as creating FCP LIFs, configuring zoning, and presenting LUNs to hosts. NetApp documents this command and its output fields in the ONTAP command reference: [system node hardware unified-connect show](#). For related FC SAN setup and verification guidance, see the [ONTAP SAN administration documentation](#).

QUESTION NO: 4

An administrator needs to configure Foreign LUN Import on a NetApp HA pair to migrate over FCP LUNs from a third-party array while serving other LUNs.

Which two actions are required to complete this task? (Choose two.)

- A. Reboot the third-party array to prepare for the LUN import.
- B. Convert all FC ports on each controller to an FC initiator port.
- C. Create zoning from the third-party array to the cluster.
- D. Convert a single FC port on each controller to an FC initiator port.

ANSWER: C D

Explanation:

Foreign LUN Import requires the ONTAP HA pair to have front-end connectivity to the foreign storage array while retaining FC target connectivity for hosts that continue to access existing NetApp LUNs. Creating zoning from the third-party array to the cluster is required so that the designated ONTAP initiator ports can discover and access the foreign LUNs. The zoning configuration must use the appropriate Fibre Channel fabric design and should expose only the intended foreign array target ports and LUNs to the NetApp initiators.

Converting a single FC port on each controller to an FC initiator port provides the HA pair with redundant import paths to the foreign array. This preserves failover capability: each node has an initiator connection for the migration workflow, while the remaining FC ports can continue operating as target ports to serve host I/O for existing NetApp LUNs. After the initiator ports and zoning are in place, ONTAP can discover the foreign LUNs and perform the supported FLI migration workflow. NetApp documents the required initiator-port configuration and SAN zoning prerequisites in its [Foreign LUN Import documentation](#) and [FC SAN configuration guidance](#).

QUESTION NO: 5

After a storage administrator creates a LUN using the CLI, the LUN does not appear on the corresponding host.

What should be done to correct this issue?

- A. Online the LUN.
- B. Map the LUN.
- C. Mount a volume.
- D. Assign an export policy.

ANSWER: B

Explanation:

Map the LUN. Creating a LUN in ONTAP creates the storage object, but it does not by itself make that LUN accessible or visible to a SAN host. To present the LUN through FCP or iSCSI, the administrator must map it to an initiator group (igroup). An igroup contains the host initiator identifiers, such as iSCSI IQNs or Fibre Channel WWPNs, that are authorized to access the LUN. The LUN mapping also assigns a LUN ID for use by the host. After the mapping is in place, the host can perform a storage rescan or rediscovery operation and detect the newly presented device. This separation between LUN creation and LUN mapping is fundamental to ONTAP SAN provisioning because it controls which hosts have access to each block device. NetApp documents LUN mapping as the operation that grants initiator access to a LUN and supports host-side discovery. See the [ONTAP SAN provisioning documentation](#) and the [lun mapping create command reference](#).

QUESTION NO: 6

An administrator connects to a LUN from Node A to a host in a redundant pathing configuration.

How does the administrator show that the LUN delivers data if a node disruption occurs?

- A. Pull all the cables from the cluster interconnect.
- B. Power off a PDU in the rack.
- C. Perform a takeover and giveback on both nodes.
- D. Shut down a backend switch.

ANSWER: C

Explanation:

Perform a takeover and giveback on both nodes. A planned takeover is the supported, controlled way to validate that an HA pair can continue serving SAN data during a node disruption. When one node is taken over, its HA partner assumes ownership of the affected node's aggregates and begins serving the LUNs hosted there. Because the host has been configured with redundant SAN paths and supported multipathing software, it can continue I/O through the available paths while storage ownership moves to the surviving controller. After confirming continued host access and I/O, giveback returns aggregate ownership to the original home node and validates the return-to-normal operation. Repeating the test for each controller validates the HA behavior regardless of which node experiences the disruption. This method tests the intended ONTAP HA mechanism while preserving a controlled operational procedure and allowing the administrator to observe host multipath behavior, path status, and application continuity. NetApp documents takeover and giveback as core storage failover operations and explains that SAN hosts require multipathing to maintain access through alternate paths. See the [ONTAP HA takeover documentation](#) and the [ONTAP SAN multipathing documentation](#).

QUESTION NO: 7

An administrator is deploying a new SAN workload that has a NetApp All SAN Array (ASA) cluster. The customer wants to reclaim space by tiering to Azure Blob storage.

According to NetApp best practices, which volume tiering policy should be applied?

- A. Snapshot-Only

- B. All
- C. None
- D. Auto

ANSWER: A

Explanation:

Snapshot-Only is the NetApp best-practice FabricPool tiering policy for a SAN workload on an ASA cluster when the goal is to reclaim performance-tier capacity. This policy keeps the volume's active file-system blocks on the local performance tier while tiering cold Snapshot copy blocks to the attached object-store tier, such as Azure Blob Storage. SAN applications commonly use latency-sensitive, random I/O patterns, so retaining the active LUN data locally provides predictable application performance. At the same time, Snapshot copies can represent a significant portion of consumed volume capacity and are generally less likely to require low-latency access. Tiering those inactive Snapshot blocks lets the customer reduce local aggregate consumption without changing how active SAN data is served. The policy is therefore well suited to a production SAN volume where Snapshot retention is needed for data protection but the primary storage tier must remain optimized for active workload I/O. NetApp's FabricPool documentation describes Snapshot-Only as tiering cold Snapshot copy data, and its tiering guidance identifies it as the appropriate policy for SAN workloads. See [ONTAP FabricPool tiering policies](#) and [ONTAP FabricPool object-store configuration](#).

QUESTION NO: 8

An engineer is planning an ONTAP upgrade for a SAN environment that includes Windows hosts, FC HBAs, and Fibre Channel switches. The engineer needs both a compatibility assessment and upgrade-specific recommendations for the existing cluster.

Which two NetApp tools should be used? (Choose two.)

- A. Interoperability Matrix Tool
- B. Active IQ Upgrade Advisor
- C. Hardware Universe
- D. NetAppDocs
- E. Harvest

ANSWER: A B

Explanation:

The Interoperability Matrix Tool is used to validate supported combinations of ONTAP releases, host operating systems, HBAs, drivers, multipathing software, and SAN switches. Active IQ Upgrade Advisor provides cluster-specific upgrade planning information and identifies risks or required actions using the system's configuration and telemetry data.

Hardware Universe is used primarily for platform hardware, adapter, cable, and port support. NetAppDocs creates implementation documentation, while Harvest is used for monitoring and metrics collection rather than upgrade compatibility analysis.

QUESTION NO: 9

An SVM is created for FCP traffic. LUNs must be created to share with ESXi hosts for datastores. Which two items must be configured after the LUN is created, for this to happen? (Choose two.)

- A. Create an igroup with the ESXi hosts ' WWPNs.
- B. Create an igroup with the ESXi hosts ' WWNNs.

C. Configure CHAP authentication.

D. Map the LUNs to the igroup.

ANSWER: A D

Explanation:

To present an FC LUN from ONTAP to ESXi for use as a datastore, an initiator group must first identify the ESXi initiators that are permitted to access the storage. For FCP, this is done by creating an igroup containing the ESXi hosts' World Wide Port Names (WWPNs). A WWPN uniquely identifies an FC host bus adapter port, allowing ONTAP to apply access control at the initiator-port level. The igroup should be created in the appropriate SVM and use the VMware host operating-system type so that the LUN presentation is aligned with the ESXi host environment.

The LUN must then be mapped to that igroup. LUN mapping creates the actual storage presentation relationship between the LUN and the authorized ESXi initiators, assigning a LUN ID through which the hosts discover and access the device over FC. Once the igroup contains the relevant ESXi WWPNs and the LUN is mapped to it, the ESXi hosts can rescan their FC adapters, detect the presented LUN, and format it as a VMFS datastore if needed. NetApp documents igroup creation and initiator identification in the [ONTAP SAN igroup documentation](#), and documents the required presentation step in [Map a LUN to an igroup](#).

QUESTION NO: 10

During maintenance of A NetApp AFF ONTAP cluster, which two steps are needed to move the SAN UFs to new nodes that are members of a portset? (Choose two.)

A. Remove the LIF from the portset.

B. Disable the SLM feature for the node that hosts the LIF.

C. Disable the zoning of the LIF WWPN.

D. Take the SAN LIF offline.

ANSWER: A D

Explanation:

To move a SAN LIF that is associated with a portset, remove the LIF from the portset before changing its node placement. A portset is a named collection of SAN LIFs that is bound to an igroup, and its membership explicitly controls the target portals that hosts can use for that igroup. Removing the LIF makes it possible to change the LIF configuration without retaining an obsolete portset association. After the LIF has been moved and is available on the intended node and port, it can be added to the appropriate portset again so that the igroup has the intended target access.

Taking the SAN LIF offline is also required because SAN LIFs are not migrated in the same manner as NAS data LIFs. The LIF must be administratively down while its home node and home port are changed. Once the new placement is configured, the LIF can be brought online and, where applicable, restored to the portset. This sequence ensures that the portset definition and the SAN LIF's physical placement remain consistent during maintenance. NetApp documents portset membership and igroup binding in its [portset configuration documentation](#) and describes SAN LIF management in the [SAN LIF management documentation](#).

QUESTION NO: 11

A customer is implementing iSCSI for business-critical LUNs on a two-node ONTAP cluster. The design must allow host I/O to continue after the failure of one Ethernet switch or one host storage NIC.

Which two design choices meet this requirement? (Choose two.)

A. Configure multiple iSCSI LIFs per node on ports connected to separate switches.

B. Configure supported host multipathing software.

- C. Connect all host and storage ports to one Ethernet switch.
- D. Configure one iSCSI LIF for the entire SVM.
- E. Apply an export policy to each LUN volume.

ANSWER: A B

Explanation:

Configuring multiple iSCSI LIFs per node across independent Ethernet switches provides redundant storage target paths. The host must also use supported multipathing software so that it can maintain and use multiple iSCSI sessions and continue I/O after a path failure.

Placing all storage connections on one switch leaves a switch as a single point of failure. A single iSCSI LIF per node does not provide the intended path redundancy, and an export policy applies to NAS access rather than iSCSI LUN access.

QUESTION NO: 12

Which two NetApp features provide synchronous data replication between two sites for SAN workloads with automatic failover in case of a site disaster? (Choose two.)

- A. SnapMirror active sync
- B. SnapMirror SVM
- C. SnapMirror Synchronous
- D. MetroCluster IP

ANSWER: A D

Explanation:

SnapMirror active sync is designed for business-continuity protection of SAN workloads across two ONTAP clusters. It synchronously replicates supported SAN LUN data and, with ONTAP Mediator, supports automated failover when a site failure occurs. This enables hosts to continue accessing the protected LUNs through the surviving cluster, while maintaining a zero recovery point objective for acknowledged writes. It is therefore directly aligned with the combination of synchronous replication, SAN support, and automatic site-failure handling described in the question.

MetroCluster IP provides an integrated synchronous replication and disaster-recovery architecture spanning two sites. It combines local high availability with replicated storage and configuration data at the partner site. In an unplanned site disaster, MetroCluster supports automated switchover when configured with the required mediator and tiebreaker infrastructure, allowing SAN workloads to be served from the surviving site. This makes it a platform-level solution for continuous availability of critical SAN environments across geographically separated locations.

NetApp documentation for [SnapMirror active sync](#) and [MetroCluster](#) describes these synchronous business-continuity capabilities.