DUMPSARENA

Implementation EngineerSAN ONTAP Exam

Netapp NS0-521

Version Demo

Total Demo Questions: 8

Total Premium Questions: 88

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, SAN Solution Design	30
Topic 2, SAN Solution Implementation	42
Topic 3, SAN Solution Validation	10
Topic 4, SAN Solution Troubleshooting	6
Total	88

QUESTION NO: 1

What is the maximum size for a LUN in NetApp ONTAP 9.14.1 software?

A. 16TiB

B. 64TiB

C. 128TiB

D. 300TiB

ANSWER: C

Explanation:

The maximum supported size of a single LUN in ONTAP 9.14.1 is **128 TiB**. A LUN is a block-storage object presented from ONTAP to SAN hosts through protocols such as FC, iSCSI, or NVMe/FC where applicable. This limit applies to the individual LUN object, rather than to the total capacity that can be provisioned across a storage VM, volume, aggregate, or cluster. Administrators designing large SAN deployments should therefore use multiple LUNs when an application requires block capacity beyond 128 TiB, while also considering host operating-system, multipathing, filesystem, and application limits. The limit supports very large enterprise database, virtualization, and application workloads while maintaining ONTAP's supported SAN configuration boundaries. NetApp's ONTAP SAN documentation describes LUN provisioning and management concepts, and NetApp maintains platform and ONTAP release-specific supported limits in its documentation and interoperability resources. See the ONTAP SAN administration documentation and the ONTAP documentation portal for current release-specific guidance.

QUESTION NO: 2

A customer has data in a LUN on an existing third-party SAN storage system that will be migrated to a new NetApp AFF A400 all-flash MetroCluster system. The host connected to the LUN is a Windows Server 2019 system.

In this scenario, which two methods are supported? (Choose two.)

A. SnapMirror Business Continuity

B. Foreign LUN Import in online mode

C. Foreign LUN Import in offline mode

D. Client-side copy of the data

ANSWER: B C

Explanation:

Foreign LUN Import in online mode and Foreign LUN Import in offline mode are supported approaches for migrating a third-party SAN LUN to ONTAP storage. Foreign LUN Import is designed specifically for block-level migration from qualified external arrays into a destination NetApp LUN. It supports Windows hosts and preserves the LUN contents and host-visible identity characteristics required to transition the host to the NetApp destination. Online import enables a controlled, minimally disruptive migration: after the initial copy, host I/O can continue while ONTAP imports changed blocks before the final cutover. Offline import is appropriate when the application can tolerate an outage; host access is stopped during the import process, simplifying consistency management for the migrated LUN. The AFF A400 platform is supported as an ONTAP SAN destination, and MetroCluster provides the destination storage system's site-to-site availability architecture rather than changing the fundamental host operating-system requirements for the import workflow. Before performing either workflow, validate the precise ONTAP release, external array model, SAN connectivity, and multipathing configuration against the Interoperability Matrix and the Foreign LUN Import documentation. See the NetApp Foreign LUN Import documentation and the NetApp Interoperability Matrix Tool.

QUESTION NO: 3

When using tagged VLANs on Cisco Nexus switches for NVMe over TCP, which two changes must be made to enable jumbo frames? (Choose two.)

A. Modify the Cisco Nexus switches to use an MTU of 9216.

B. Modify the appropriate broadcast domain in NetApp ONTAP software to use an MTU of 9000.

C. Modify the appropriate broadcast domain in NetApp ONTAP software to use an MTU of 9216.

D. Modify the Cisco Nexus switches to use an MTU of 9000.

ANSWER: A B

Explanation:

Jumbo-frame support for NVMe/TCP must be configured consistently across the end-to-end data path, but the MTU values are expressed differently by Cisco Nexus switching and ONTAP. The Cisco Nexus switches must use an MTU of 9216. Nexus platforms commonly use 9216 bytes as the Layer 2 system or network QoS MTU so that a 9000-byte IP payload can traverse the switch while retaining room for Ethernet headers and VLAN tagging. This setting must apply to every relevant switch interface and path used by the tagged VLAN.

The ONTAP broadcast domain must use an MTU of 9000. ONTAP's broadcast-domain MTU represents the IP network MTU configured for the associated ports, VLANs, and IP interfaces. A 9000-byte ONTAP MTU aligns with the intended jumbo IP packet size for NVMe/TCP, while the Nexus 9216-byte setting accommodates the encapsulation overhead at the switching layer. Configuring these values provides the required end-to-end jumbo-frame capability for the NVMe/TCP network. NetApp documents the 9000-byte ONTAP network MTU requirement for NVMe/TCP configuration, and Cisco documents the Nexus QoS MTU behavior and configuration. See NetApp ONTAP NVMe/TCP configuration documentation and Cisco Nexus 9000 QoS configuration guide.

QUESTION NO: 4

A customer asks for help cloning one database LUN from a SnapMIrror destination volume for testing the data with a new application.

What command should the administrator use to clone the LUN?

A. volume file clone create

B. lun move start

C. volume clone create

D. volume move start

ANSWER: C

Explanation:

`volume clone create` is the appropriate command because the source is a SnapMirror destination volume and the requirement is to make its database LUN available safely for application testing. This command creates a writable FlexClone volume using the SnapMirror destination volume, or a Snapshot within it, as its parent. The resulting FlexClone is space-efficient because it initially shares blocks with the parent rather than requiring a full physical copy. The cloned volume contains the database LUN and its surrounding volume content in a consistent, independently writable test environment. This is particularly useful when testing a new application against replicated production data: the SnapMirror destination remains the protected replication copy, while changes made during testing occur in the FlexClone volume. The clone can be mounted or mapped to the testing hosts after the appropriate SAN configuration is completed. NetApp documents FlexClone creation from SnapMirror destination volumes as a supported workflow, including creation of a writable clone from destination data. See the NetApp guidance for creating FlexClone volumes from SnapMirror destinations and the volume clone create command reference.

QUESTION NO: 5

A customer has created an SVM for their SAN workloads. They now want to configure the SVM to use NVMe/FC. Which two steps are needed to accomplish this task? (Choose two.)

A. Add the FC protocol.

B. Create the FC service.

C. Create an NVMe/FC LIF.

D. Create the NVMe service.

ANSWER: C D

Explanation:

Creating the NVMe service is required because it enables the SVM to provide NVMe storage services. ONTAP uses the NVMe service to expose NVMe namespaces and subsystems from that SVM; this is separate from the FCP service used for traditional SCSI Fibre Channel workloads. After the NVMe service is enabled, an NVMe/FC LIF provides the FC network endpoint through which NVMe hosts connect to the SVM. The LIF is configured with the NVMe/FC data protocol and must be placed on a suitable FC-capable node and port so that host initiators can discover and access the NVMe subsystem. Together, the NVMe service supplies the NVMe control-plane capability for the SVM, while the NVMe/FC LIF supplies the required data-path connectivity over Fibre Channel. NetApp's NVMe over Fibre Channel configuration workflow specifically includes creating the NVMe service and creating NVMe/FC LIFs. See the ONTAP NVMe over Fibre Channel configuration documentation and the procedure for creating NVMe/FC LIFs.

QUESTION NO: 6

An administrator creates a new volume and a LUN for an SAN-connected application that uses snapshots. The application has fast-growing data that is constantly being populated and deleted afterwards.

Which option can be set to manage the space automatically?

A. Enable automatic resizing on the volume.

B. Enable space reservation on the LUN.

C. Disable volume snapshot autodelete.

D. Unmap the LUN.

ANSWER: A

Explanation:

Enable automatic resizing on the volume. ONTAP volume autosize provides automatic capacity management by increasing a volume's size when its used-space threshold is reached and, when configured for grow-and-shrink behavior, reducing the volume size again as space use falls. This is appropriate for an application with rapidly changing data consumption because it reduces the need for an administrator to continually adjust the volume's provisioned size manually. The LUN remains hosted in the same FlexVol volume while ONTAP manages the volume boundary, provided that sufficient aggregate capacity is available for growth. In a SAN environment, the underlying volume must have enough available space to support LUN writes and Snapshot-related block retention. Autosize settings, including the maximum permitted volume size and growth increment, should therefore be selected according to expected workload peaks and aggregate capacity. NetApp documents volume autosize as a feature for automatically growing or shrinking FlexVol volumes in response to space utilization; see Enable automatic growth of a volume and ONTAP FlexVol volume management commands.

QUESTION NO: 7

A storage engineer in SAN environment needs to Identify FC ports using the company-wide naming policy. Which command will accomplish this task?

A. vserver fcp wwpn-alias set

B. network interface modify

C. TCP portname set

D. ucadmin modify

ANSWER: A

Explanation:

The `vserver fcp wwpn-alias set` command is correct because it assigns a human-readable alias to a Fibre Channel World Wide Port Name (WWPN) within an SVM. A WWPN is the globally unique hardware identifier used for an FC port, but its hexadecimal format is not convenient for operational use. Defining an alias lets the storage team apply a consistent company naming convention, such as identifying the host, adapter, fabric, or port role associated with that WWPN. These aliases make FC port identities easier to recognize when administering SAN connectivity and reviewing Fibre Channel configuration information, without changing the underlying WWPN or the FC port's physical identity. The command is used in the context of the target SVM and associates the selected WWPN with the chosen descriptive alias. This directly supports the requirement to identify FC ports according to an organization-wide naming policy. NetApp documents the command and its parameters in the ONTAP CLI command reference. Additional FC protocol administration context is available in the ONTAP SAN administration documentation.

QUESTION NO: 8

A customer is implementing SnapMirror active sync for a business-critical FC workload between two ONTAP clusters. Cluster peering is already complete. The solution must support Automated Unplanned Failover when a protected site becomes unavailable.

Which two additional configurations are required? (Choose two.)

A. Configure an SVM peer relationship.

B. Configure ONTAP Mediator connectivity.

C. Configure an FC interswitch link between the sites.

D. Bind the FC igroup to an iSCSI portset.

E. Apply a SnapVault policy to the protected LUNs.

ANSWER: A B

Explanation:

An SVM peer relationship is required for the participating SVMs to support the intercluster data-protection relationship. ONTAP Mediator must also be configured and reachable so that SnapMirror active sync can use quorum information for Automated Unplanned Failover decisions.

An FC interswitch link is part of a fabric design and is not a SnapMirror active sync requirement between the ONTAP clusters. A portset controls iSCSI target LIF access, not synchronous replication behavior. A SnapVault policy does not provide the active-active synchronous protection and automated failover behavior required here.