

DUMPSBOSS.

Network Appliance Data Protection Solutions

Netapp NS0-163

Version Demo

Total Demo Questions: 20

Total Premium Questions: 205

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

Which two networking components must be configured to failover LIFs across subnets? (Choose two.)

- A. BGP
- B. LACP
- C. VLAN
- D. VIP

ANSWER: A D

Explanation:

BGP and **VIP** are required for this use case. In ONTAP, a virtual IP (VIP) LIF provides an IP address that is not tied to one physical port or to a single local subnet in the way that a conventional data LIF is. The VIP address can be hosted by an eligible node and can move when a node or network path becomes unavailable. This supports highly available client access across the network rather than limiting access to a local broadcast domain.

Border Gateway Protocol (BGP) supplies the routing component that makes the VIP LIF reachable. ONTAP establishes BGP peering with upstream routers and advertises routes for the VIP address from the node currently hosting it. When the VIP LIF fails over, ONTAP withdraws and re-advertises the route as appropriate, allowing routers to direct client traffic to the new hosting location. Thus, VIP provides the movable endpoint and BGP distributes the dynamic routing information needed for cross-subnet reachability and failover. NetApp documents VIP LIF configuration and its BGP requirements in the [ONTAP VIP LIF documentation](#) and the [ONTAP BGP configuration documentation](#).

QUESTION NO: 2

Which one of the following commands is used to specify the destination appliance(s) that are allowed to replicate (pull) data from the source appliance?

- A. options snapmirror.allow
- B. snapmirror release
- C. snapmirror initialize
- D. options snapmirror.access

ANSWER: D

Explanation:

`options snapmirror.access` is the Data ONTAP 7-Mode configuration command used on the source appliance to define which remote appliances are authorized to access it for SnapMirror replication. The value can identify one or more permitted destination systems, typically by host name or IP address. This source-side authorization is important for pull-based replication because the destination appliance contacts the source to read Snapshot copy data and transfer it into the destination volume. Configuring `snapmirror.access` therefore establishes the set of destination appliances that may initiate and perform SnapMirror data transfers from that source. In the legacy 7-Mode command structure reflected by this question, the setting is managed through the `options` command rather than as a replication relationship operation. Current ONTAP documentation continues to describe SnapMirror as a source-to-destination replication relationship and documents the administrative model for creating and managing those relationships. See the [NetApp ONTAP documentation for creating replication relationships](#) and the [NetApp ONTAP CLI reference for snapmirror create](#).

QUESTION NO: 3

Which are three valid data protection solutions? (Choose three.)

- A. SnapVault to Volume SnapMirror
- B. SnapVault to Volume SnapMirror to SnapVault
- C. SnapVault to Volume SnapMirror to Volume SnapMirror
- D. Volume SnapMirror to SnapVault

ANSWER: A C D

Explanation:

NetApp ONTAP supports cascaded data-protection topologies in which a replicated destination volume can serve as the source for a subsequent Volume SnapMirror relationship. Therefore, **SnapVault to Volume SnapMirror** is valid when a vault copy is used as the source of a downstream mirror. This can provide a recoverable protected copy at the intermediate site while maintaining a further replicated copy elsewhere.

SnapVault to Volume SnapMirror to Volume SnapMirror is also valid. It extends that topology by using the first downstream Volume SnapMirror destination as the source for another Volume SnapMirror relationship, creating a multi-site cascade. Each hop uses compatible data-protection relationships and preserves the intended replication flow.

Volume SnapMirror to SnapVault is valid because a mirrored destination can be used as the source for a SnapVault relationship. This arrangement combines rapid disaster-recovery replication through Volume SnapMirror with longer-term, policy-based snapshot retention through SnapVault at the downstream destination. In practice, administrators configure relationship types and policies to meet recovery-point, retention, and site-protection requirements. NetApp documents replication relationships and supported protection configurations in the [ONTAP replication relationship documentation](#) and the [ONTAP data protection documentation](#).

QUESTION NO: 4

A SnapMirror destination was broken during a disaster, and applications have been writing to it for several days. The original source volume is available again and must become the new destination without discarding the application changes made during the outage.

Which resynchronization direction should the administrator use?

- A. Resynchronize from the former source to the current active volume.
- B. Resynchronize from the current active volume to the former source.
- C. Initialize the former source from the current active volume.
- D. Quiesce the relationship on the current active volume.

ANSWER: B

Explanation:

The administrator should resynchronize from the current active volume to the former source volume. Because the broken destination accepted new writes during the outage, it contains the authoritative data set. Resynchronizing in this direction makes the former source a destination and preserves the changes made while clients used the disaster-recovery site.

Resynchronizing from the former source to the active volume could overwrite the newer data written during the outage. Initializing the relationship would require a new baseline transfer and is unnecessary when an appropriate common Snapshot copy is available.

QUESTION NO: 5

Which ONTAP policy type defines the kind of traffic that is supported for a LIF?

- A. QoS

- B. Service
- C. Failover
- D. Export

ANSWER: B

Explanation:

Service is correct because an ONTAP LIF service policy determines the services, and therefore the traffic types, that are allowed on a logical interface. A service policy is assigned when a LIF is created or modified and contains one or more predefined services, such as data NFS, data SMB, iSCSI, management HTTPS, or intercluster core. This association defines the LIF's intended function and helps ONTAP ensure that the LIF is used for the appropriate protocol or management traffic. For example, a LIF configured with data services supports client data access for the applicable NAS or SAN protocol, while a cluster-management-oriented policy enables administrative traffic. ONTAP uses the service-policy model to provide a consistent, role-based way to configure LIF capabilities across supported ONTAP releases. Administrators can view service policies and their included services with ONTAP CLI commands such as `network interface service-policy show`, and can assign a policy through LIF configuration. NetApp documents that service policies define the network services supported by a LIF and replace the earlier LIF role-based approach. See the [NetApp ONTAP documentation on supported LIF traffic](#) and the [LIF configuration overview](#).

QUESTION NO: 6

What are two logical network components of ONTAP? (Choose two.)

- A. VLAN
- B. SFP+
- C. Network port
- D. Interface group

ANSWER: A D

Explanation:

In ONTAP, a VLAN and an interface group are logical network components that are configured on top of physical network ports. A VLAN logically separates network traffic by associating a port with an IEEE 802.1Q VLAN tag. ONTAP represents this configuration as a logical VLAN port, enabling storage traffic to use the appropriate Layer 2 network segment without requiring a dedicated physical adapter for every segment. An interface group, also called an ifgrp, logically combines two or more physical ports into one aggregated logical interface. This provides increased bandwidth and, depending on the selected distribution function and switch configuration, resiliency through multiple member links. Logical ports such as VLANs and interface groups can then be used in ONTAP networking configurations, including assignment to broadcast domains and hosting LIFs. NetApp's networking documentation distinguishes these logical port types from the underlying physical ports on a controller. See the [ONTAP physical network ports documentation](#) and the [ONTAP VLAN configuration documentation](#).

QUESTION NO: 7

Which three tools present statistics from the Data ONTAP counter manager? (Choose three.)

- A. stats
- B. sysstat
- C. windows perfmon
- D. operations manager

ANSWER: A C D

Explanation:

stats, **windows perfmon**, and **operations manager** can present performance statistics collected by the Data ONTAP counter manager. The **stats** command is the native Data ONTAP command-line interface for identifying counter objects, instances, and counters, and for displaying counter values directly from the cluster. This makes it the primary low-level interface for examining ONTAP performance data. Windows Performance Monitor can display Data ONTAP performance counters when the applicable ONTAP performance-counter integration is configured, allowing administrators to graph and analyze storage statistics alongside Windows performance data. Operations Manager consumes Data ONTAP monitoring and performance information and presents it in a centralized management interface, supporting historical analysis, reporting, and operational monitoring. These interfaces all use the counter-based performance-statistics model rather than relying solely on a point-in-time system summary. NetApp's current CLI documentation describes viewing performance-counter data through the statistics interfaces: [ONTAP statistics show command reference](#). NetApp also documents performance monitoring and counter collection concepts in its ONTAP administration documentation: [ONTAP performance monitoring overview](#).

QUESTION NO: 8

Which two configurations ensure end users are able to restore previous versions of files? (Choose two.)

- A. SVM previous versions enabled
- B. Volume snapshot copy
- C. Enable FlexClone operations
- D. Share previous versions enabled

ANSWER: B D

Explanation:

End-user restoration through the Windows *Previous Versions* interface depends on ONTAP Snapshot copies and on exposing those copies through the SMB share. A **Volume snapshot copy** provides the point-in-time, read-only version of the volume from which users can browse and recover an earlier file or folder version. Snapshot copies are space-efficient and are the underlying ONTAP technology used to present prior versions of SMB-accessible data.

Share previous versions enabled permits the SMB share to display available Snapshot copies to Windows clients as previous versions. In ONTAP, this behavior is controlled by the SMB share `show-previous-versions` setting. When it is enabled and the hosting volume contains eligible Snapshot copies, users can open a file or folder's Properties dialog in Windows, select the Previous Versions tab, and copy or restore a prior version without requiring administrator intervention. This capability combines volume-level Snapshot retention with the share-level presentation setting, providing self-service recovery while retaining ONTAP Snapshot efficiency. See NetApp's [documentation on restoring previous SMB file versions](#) and the [vserver cifs share create command reference](#).

QUESTION NO: 9

An administrator wants to scale out storage performance for real time read-intensive workloads across multiple controllers.

Which two NetApp ONTAP features accomplish this task? (Choose two.)

- A. FlexCache
- B. FlexGroup
- C. FlexVol
- D. FlexClone

ANSWER: A B

Explanation:

FlexCache and **FlexGroup** provide ONTAP scale-out capabilities for read-intensive workloads across controllers. FlexCache creates cache volumes that serve data close to clients while retaining an authoritative origin volume. Client read requests can be distributed to cache volumes hosted on different nodes or clusters, reducing load and latency at the origin and allowing read throughput to scale horizontally. This is particularly useful when many clients repeatedly access the same data set, including geographically distributed or high-performance computing workloads.

FlexGroup provides a single namespace that is transparently composed of multiple constituent volumes. Those constituents can be placed across nodes in an ONTAP cluster, enabling aggregate throughput and capacity from multiple controllers while clients continue to access one logical volume. ONTAP distributes files among the constituents, allowing parallel client I/O and avoiding a single-volume performance bottleneck. Together, these features address read scale-out at different layers: FlexCache scales reads through distributed caching, and FlexGroup scales a namespace through clustered constituents. NetApp documents FlexCache read scaling at [ONTAP FlexCache documentation](#) and describes FlexGroup scale-out volume architecture at [ONTAP FlexGroup documentation](#).

QUESTION NO: 10

A user accidentally deletes a single file from an SMB share. A recent Snapshot copy contains the required file, and the administrator must avoid reverting changes made to other files in the volume after that Snapshot copy was created.

Which recovery method should the administrator use?

- A. Perform a volume SnapRestore operation
- B. Perform a single-file SnapRestore operation
- C. Break the SnapMirror relationship
- D. Create a FlexClone volume

ANSWER: B

Explanation:

Single-file SnapRestore is the appropriate recovery method. It restores the affected file from a selected Snapshot copy while leaving the remainder of the active volume unchanged. This provides the required granular recovery without discarding subsequent changes to unrelated files.

A volume SnapRestore operation reverts the complete volume to the selected Snapshot copy and would remove later changes across the volume. Breaking a SnapMirror relationship is a disaster-recovery operation and does not restore an individual file. Creating a FlexClone can provide a writable copy for investigation, but it does not itself restore the deleted production file.

QUESTION NO: 11

Which two products can be used to administer CloudVolumes ONTAP? (Choose two.)

- A. Cloud Data Sense
- B. System Manager
- C. NetApp BlueXP
- D. CloudCheckr

ANSWER: B C

Explanation:

Cloud Volumes ONTAP can be administered through **System Manager** and **NetApp BlueXP**. BlueXP provides the cloud-based management plane for deploying and managing Cloud Volumes ONTAP working environments across supported cloud providers. From the BlueXP console, administrators can perform operational tasks such as creating and managing working environments, configuring capacity and storage services, monitoring resources, and integrating data-protection services. NetApp documentation identifies the BlueXP working environment as the central interface for Cloud Volumes ONTAP administration.

System Manager is the ONTAP web-based management interface and can also manage a Cloud Volumes ONTAP cluster directly. It is appropriate for ONTAP administrative activities, including configuring storage virtual machines, volumes, networking, data protection, users, and other cluster-level settings. This gives administrators both a cloud-management experience through BlueXP and a familiar ONTAP-native interface through System Manager. NetApp's Cloud Volumes ONTAP documentation describes management through BlueXP, while ONTAP documentation covers System Manager access and administration. See [Cloud Volumes ONTAP documentation](#) and [ONTAP System Manager documentation](#).

QUESTION NO: 12

Which one of the following statements apply to Advanced Single Instance Storage (ASIS)?

- A. Eliminates duplicates files in SnapVault for NetBackup secondary volumes
 - B. Eliminates duplicates blocks in SnapVault for NetBackup source volumes
 - C. Eliminates duplicates blocks in SnapVault for NetBackup secondary volumes
 - D. Eliminates duplicates files in SnapVault for NetBackup source volumes
- Correct answers

ANSWER: C**Explanation:**

Eliminates duplicates blocks in SnapVault for NetBackup secondary volumes is correct. Advanced Single Instance Storage (ASIS) is NetApp's block-level deduplication capability designed for SnapVault for NetBackup environments. It operates on the SnapVault secondary system, where protected backup images are retained, identifying identical data blocks and storing only one physical instance while maintaining the logical references needed for each backup image. This reduces the disk capacity required on the backup target without changing the way NetBackup accesses or catalogs the protected data. The focus on blocks is important: ASIS recognizes duplicate content at a granularity below whole files, allowing storage savings even when files are not identical but contain common blocks. Applying this efficiency on the secondary volume also aligns with the purpose of the SnapVault for NetBackup architecture: preserve backup data on a dedicated NetApp destination while improving backup-storage utilization. NetApp describes deduplication as a storage-efficiency technology that eliminates redundant data blocks, and its data-protection documentation explains the role of destination volumes in SnapVault replication. See [NetApp ONTAP storage efficiency concepts](#) and [NetApp data protection replication workflow](#).

QUESTION NO: 13

Which two statements apply to a FlexClone volume created from a Snapshot copy? (Choose two.)

- A. It is created as a read-only copy of the parent volume.
- B. It is writable and initially shares blocks with the parent data.
- C. Changes written to the clone are automatically copied to the parent volume.
- D. Its backing Snapshot copy is retained while the clone depends on it.

ANSWER: B D**Explanation:**

A FlexClone volume is a writable clone that is created rapidly from a parent volume or a Snapshot copy. It is initially space-efficient because it shares data blocks with its parent rather than copying the entire data set. As changes are made to the clone or parent, ONTAP stores changed blocks separately.

The parent Snapshot copy on which the clone depends cannot be deleted while the clone still requires it. A FlexClone volume is not an independent full copy at creation time, and it does not automatically replicate its changes back to the parent volume.

QUESTION NO: 14

A production site is unavailable, and the SnapMirror destination volume contains the latest replicated data. The administrator must make the destination volume writable so that clients can resume access during the outage.

Which action should the administrator perform?

- A. Quiesce the SnapMirror relationship
- B. Update the SnapMirror relationship
- C. Break the SnapMirror relationship
- D. Resynchronize the SnapMirror relationship

ANSWER: C

Explanation:

SnapMirror break is correct. Breaking the SnapMirror relationship changes the destination volume from read-only to read-write, allowing it to serve client I/O during a disaster-recovery event. The destination retains the replicated contents and Snapshot copies that were present when the relationship was broken.

`snapmirror quiesce` stops new transfers but does not make the destination writable. `snapmirror update` performs an incremental transfer, and `snapmirror resync` is used to re-establish a replication relationship after the roles or contents have diverged.

QUESTION NO: 15

Which OSSV package would you use to perform an OSSV unattended installation?

- A. `svunattendossv`
- B. `unattendpackager`
- C. `svconfigpackager`
- D. None of the above

ANSWER: C

Explanation:

svconfigpackager is the OSSV (SnapVault Open Systems) utility used to create an unattended-installation package. It captures the installation and configuration choices needed for a consistent OSSV deployment, allowing the same package to be used when installing OSSV on additional supported hosts without stepping through the interactive installer each time. This is particularly useful for standardized rollouts, scripted server provisioning, and environments that require repeatable configuration of protection software. The package is generated before deployment and contains the inputs that the OSSV installer needs in order to run without manual intervention. OSSV is a legacy NetApp product, so its installation materials and downloads are generally accessed through the NetApp Support site rather than current ONTAP documentation. Refer to the SnapVault Open Systems product download area on the [NetApp Support Site](#) and NetApp's current [data-protection documentation](#) for supported data-protection product guidance.

QUESTION NO: 16

Fill in the blank for the following statement using the best answer from the choices that follow. A snapshot is a frozen, read-only image of the entire Data ONTAP file system that reflects the state of the _____ at the time the snapshot was created.

- A. C
- B. file system

ANSWER: B

Explanation:

The correct completion is **file system**. In traditional Data ONTAP terminology, a Snapshot copy is a frozen, read-only image that preserves the state of a file system at the instant the Snapshot copy is created. It provides a point-in-time representation without requiring a full physical copy of the data. ONTAP accomplishes this through its write-anywhere file layout and copy-on-write behavior: blocks referenced by the Snapshot remain available in their original form while later changes are written elsewhere. Consequently, users and data-protection operations can access the consistent file-system state represented by that Snapshot even after active data has changed. This point-in-time preservation is the basis for restoring deleted or changed files and for using Snapshot copies as a source for replication and backup workflows. Modern ONTAP documentation commonly describes the same construct in volume-oriented terms, because Snapshot copies are created for volumes, but the wording in this question specifically uses the classic definition of the Data ONTAP file system. See NetApp's [documentation on local Snapshot copies](#) and its [Snapshot copy concepts](#) for the point-in-time behavior and underlying storage efficiency.

QUESTION NO: 17

Which two of the following statements apply to the `fs_size_fixed` volume option?

- A. The default setting is on
- B. This option is automatically set to be on when a volume becomes a Snap Mirrored volume
- C. This option causes the file system to remain the same size and not grow or shrink when a Snap Mirrored volume relationship is broken
- D. The option turns to off after the `snapmirror break` command is issued for the volume

ANSWER: B C

Explanation:

The `fs_size_fixed` setting controls whether a volume's file system automatically changes size when the containing volume size changes. NetApp automatically enables this setting when a volume is made a SnapMirror destination, ensuring that the destination file system maintains the size required by replicated data rather than being automatically resized. This behavior is particularly important for a data-protection destination because its contents and sizing are governed by replication from the source volume.

When the SnapMirror relationship is broken, the volume becomes writable, but the fixed-file-system setting remains enabled. Consequently, the file system remains at its established size and does not automatically grow or shrink as a result of the break operation. An administrator can subsequently change the volume configuration if a different behavior is required. NetApp documents `-fs-size-fixed` as the volume-modify parameter that determines whether the file system is fixed, and documents `SnapMirror break` as the operation that makes a destination volume writable while retaining its volume configuration. See the [ONTAP volume modify command reference](#) and the [ONTAP snapmirror break command reference](#).

QUESTION NO: 18

What are two properties of the NVMe protocol? (Choose two.)

- A. LUNs are mapped to Initiator groups.
- B. Namespaces are mapped to subsystems.
- C. NQNs are added to namespaces.
- D. NQNs are added to subsystems.

ANSWER: B D

Explanation:

In ONTAP NVMe, a namespace is the storage object presented to an NVMe host, analogous in purpose to a block device. To make that namespace available, ONTAP maps it to an NVMe subsystem. The subsystem acts as the access-control and connectivity container through which hosts can discover and access the mapped namespaces. NVMe Qualified Names (NQNs) uniquely identify NVMe entities, including hosts and subsystems. ONTAP associates an NVMe host NQN with an NVMe subsystem, thereby defining which host is permitted to access the namespaces mapped to that subsystem. Consequently, "Namespaces are mapped to subsystems." and "NQNs are added to subsystems." describe the relevant NVMe/ONTAP relationships. NetApp's ONTAP documentation describes creating NVMe subsystems, adding host NQNs to them, and mapping namespaces to those subsystems as the required configuration sequence for NVMe host access. See [NetApp ONTAP NVMe configuration overview](#) and [NetApp documentation for NVMe namespaces and subsystems](#).

QUESTION NO: 19

An administrator enables the NFS protocol on an SVM and creates a volume to be used as a datastore in a VMware environment. After a rescan, the volume is not visible on the ESXi hosts.

What must the administrator validate to resolve this problem?

- A. Export policy
- B. Name service
- C. Security style
- D. Access control

ANSWER: A

Explanation:

Export policy is correct because ONTAP controls NFS client access to a volume through the export policy associated with that volume or its containing junction path. Creating a volume and enabling NFS on the SVM does not, by itself, permit an ESXi host to mount the volume as an NFS datastore. The administrator must verify that the applicable export policy contains an export rule that matches the ESXi host management/IP network and grants the required NFS access, normally read-write access for a datastore. The rule's client match, access protocol, read-write permissions, and superuser handling must be appropriate for the ESXi hosts and the NFS version in use. The volume must also be mounted at its configured junction path, because that path is what ESXi uses when mounting the datastore. Once the export policy allows the ESXi hosts, the hosts can mount or discover the NFS datastore during the rescan/add-datastore operation. NetApp documents that export policies and their rules define which NFS clients can access data and with what permissions. See [NetApp ONTAP: Add a rule to an export policy](#) and [NetApp ONTAP: Create and configure an NFS volume](#).

QUESTION NO: 20

You cannot disable the SyncMirror license if one or more mirrored volumes exist.

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. SyncMirror is a licensed ONTAP data-availability feature that maintains mirrored storage by using separate plexes. A volume that is configured as mirrored depends on the SyncMirror capability remaining licensed and available. Consequently, ONTAP protects the consistency and manageability of that configuration by preventing removal or disabling of the SyncMirror license while one or more mirrored volumes still exist. Before the license can be removed, the administrator must first remove the SyncMirror relationship from every affected mirrored volume, ensuring that no active volume configuration relies on the feature. This dependency check avoids leaving volumes in a state where their configured storage layout requires functionality that is no longer licensed. NetApp's licensing documentation describes how installed licenses are managed and how feature usage can affect license removal, while its storage-management documentation covers SyncMirror's mirrored-plex design and administration. See [NetApp ONTAP license management documentation](#) and [NetApp documentation for managing plexes and SyncMirror](#).