

DUMPSBOSS.

Hybrid Cloud Administrator Exam

Netapp NS0-304

Version Demo

Total Demo Questions: 8

Total Premium Questions: 82

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Manager	19
Topic 2, Cloud Volumes ONTAP	35
Topic 3, Cloud Backup Service	5
Topic 4, Cloud Sync	3
Topic 5, Cloud Tiering	6
Topic 6, Mix Questions	14
Total	82

QUESTION NO: 1

An administrator is asked to set up a Cloud Volumes ONTAP (CVO) with high availability in AWS using all default configuration settings. Where is the IAM role created?

- A. Cloud Volumes ONTAP
- B. BlueXP
- C. AWS Systems Manager
- D. AWS console

ANSWER: B

Explanation:

BlueXP is correct. When an administrator deploys an AWS high-availability Cloud Volumes ONTAP working environment with the default settings, BlueXP performs the deployment orchestration through its Connector. As part of that automated process, BlueXP uses the AWS permissions granted to the Connector to create the required Cloud Volumes ONTAP IAM role and instance profile in the selected AWS account. It then attaches the instance profile to the Cloud Volumes ONTAP instances so that they can securely make the AWS API calls needed for normal operation, such as storage, networking, and high-availability functions.

This behavior depends on the Connector having the NetApp-required AWS permissions, including permissions to create, manage, and pass IAM roles and instance profiles. The administrator therefore does not need to manually create the Cloud Volumes ONTAP node IAM role in the AWS console when using the standard BlueXP deployment workflow and its defaults. BlueXP automates the AWS resource creation while AWS IAM remains the underlying service where the resulting role is stored and enforced. NetApp documents the required AWS permissions and the Cloud Volumes ONTAP deployment process in the [AWS permissions reference](#) and the [Cloud Volumes ONTAP AWS deployment documentation](#).

QUESTION NO: 2

An administrator notices that Cloud Data Sense is not scanning the new NFS volume that was recently provisioned. What should the administrator enable?

- A. S3 access
- B. Read permissions
- C. CIFS access
- D. Write permissions

ANSWER: B

Explanation:

Read permissions must be enabled. BlueXP classification, formerly called Cloud Data Sense, scans NFS data by mounting the exported volume and reading its directory and file contents. For a newly provisioned volume, the administrator must ensure that the volume's export policy contains a rule allowing the BlueXP classification instance or scanner host to access the export with read-only or read-write permissions. The rule must apply to the scanner's source IP address or network and use the applicable NFS protocol version and security type. Read-only access is sufficient for classification because the service needs to discover files and inspect their contents and metadata; it does not need to modify production data to perform a standard scan. After the export-policy rule is in place and the volume is selected in the BlueXP classification configuration, the service can mount the NFS volume and begin or resume scanning it. NetApp's setup guidance specifically identifies NFS export-policy access for the classification instance as a required prerequisite for scanning ONTAP NFS volumes. See [BlueXP classification scanning prerequisites for ONTAP volumes](#) and [ONTAP NFS export-policy configuration](#).

QUESTION NO: 3

An administrator has an Azure NetApp Files volume in a Standard capacity pool. The workload now requires the throughput characteristics of a Premium service level, and the administrator must preserve the existing volume data.

What should the administrator do?

- A. Increase the volume quota in the existing capacity pool.
- B. Move the volume to a Premium capacity pool.
- C. Create a new Premium volume and recreate the data manually.
- D. Change the protocol type of the existing volume.

ANSWER: B

Explanation:

An Azure NetApp Files volume inherits its service level from the capacity pool in which it resides. To change from Standard to Premium while preserving the volume and its data, the administrator should move the volume to a capacity pool that uses the Premium service level.

Changing the volume quota changes allocated capacity, not its service level. Creating a new volume would require a separate data-migration process and is unnecessary when an appropriate destination capacity pool is available.

QUESTION NO: 4

An administrator is deploying FlexCache volumes between a Production SVM and a Development SVM on the same 8-node cluster. Which network is being used?

- A. NAS data LIFs
- B. Node Management
- C. IntraCluster
- D. InterCluster

ANSWER: C

Explanation:

IntraCluster is the network used when the FlexCache origin volume and FlexCache volume are hosted in separate SVMs on the same ONTAP cluster. In this topology, ONTAP communicates internally between the cache and origin through the cluster network, which provides the back-end connectivity among nodes in a cluster. This communication supports cache fills, validation, metadata operations, and requests that must be serviced by the origin volume.

FlexCache is designed to provide scalable, distributed access to NAS data while retaining a single authoritative origin volume. When both the origin and cache are local to one cluster, ONTAP uses intra-cluster connectivity rather than requiring intercluster LIFs. Intercluster networking is relevant when the origin and cache reside on different clusters. The fact that the Production SVM and Development SVM are distinct does not change this behavior: their placement on the same eight-node cluster makes this an intra-cluster FlexCache relationship.

NetApp documents that FlexCache can be deployed within a cluster or across clusters, with the networking path determined by that placement. See the [ONTAP FlexCache documentation](#) and the [ONTAP network management reference](#).

QUESTION NO: 5

An administrator is planning an AWS Cloud Volumes ONTAP HA deployment that must tolerate an Availability Zone failure.

Which two placement requirements should the administrator use? (Choose two.)

- A. Deploy the two Cloud Volumes ONTAP nodes in separate Availability Zones.
- B. Deploy the mediator in a separate Availability Zone.
- C. Deploy both Cloud Volumes ONTAP nodes in the same Availability Zone.
- D. Deploy all HA components in public subnets.
- E. Deploy an AWS Transit Gateway in the mediator subnet.

ANSWER: A B

Explanation:

The two Cloud Volumes ONTAP HA nodes must be deployed in separate Availability Zones so that an Availability Zone failure does not affect both storage nodes. The mediator is deployed in a separate Availability Zone to provide the independent arbitration component used by the HA configuration.

Placing both nodes in one Availability Zone creates a single Availability Zone failure domain. A public subnet and an AWS Transit Gateway can be part of a broader network design, but neither is a required HA placement criterion.

QUESTION NO: 6

Which two options must be configured when selecting a service level for Cloud Volumes Service (CVS) in GCP? (Choose two.)

- A. Performance
- B. High-Availability
- C. Capacity
- D. RAID Level

ANSWER: A C

Explanation:

Performance and **Capacity** are configured together because they define the usable storage service delivered to the workload. Performance is selected through the service level, which determines the performance characteristics available to the volume. In Cloud Volumes Service for Google Cloud, service levels are designed for differing workload requirements, so this selection aligns the volume with the expected throughput and latency needs of applications such as file services, databases, or development workloads.

Capacity is the provisioned size of the volume. It must be set because the service needs a defined amount of storage to allocate and because the selected capacity participates in the resulting service characteristics and cost. Administrators should therefore size capacity for present data requirements while allowing suitable room for anticipated growth. Together, the performance/service-level selection and allocated capacity establish the volume configuration before clients can use the storage. Google Cloud's volume-creation guidance describes specifying both a service level and an allocated capacity as part of creating a volume: [Create a volume](#). Additional service-level details are available in the [Google Cloud NetApp Volumes overview](#).

QUESTION NO: 7

An administrator is preparing to create a SnapMirror relationship from an on-premises ONTAP source SVM to a Cloud Volumes ONTAP destination SVM.

Which two peering relationships are required before the SnapMirror relationship can be created? (Choose two.)

- A. A cluster peer relationship
- B. An aggregate peer relationship
- C. An SVM peer relationship
- D. A data LIF peer relationship
- E. A Connector peer relationship

ANSWER: A C

Explanation:

A cluster peer relationship and **an SVM peer relationship** are required. Cluster peering establishes the trusted intercluster communication path between the on-premises cluster and the Cloud Volumes ONTAP cluster. SVM peering then authorizes the participating SVMs to replicate data between one another.

An aggregate is a local storage construct and is not peered. A data LIF provides NAS or SAN client access and does not establish SnapMirror trust. A Connector manages BlueXP workflows but does not replace ONTAP cluster and SVM peering.

QUESTION NO: 8

An administrator is preparing an ONTAP NFS volume for a BlueXP Classification deep scan. The scan host can reach the SVM data LIF, but the export policy has been replaced with a restrictive rule.

Which two export-policy settings are required for the scan host? (Choose two.)

- A. Read access for the scan host
- B. Superuser access for the scan host
- C. Write access for the scan host
- D. Kerberos-only authentication for the scan host
- E. An anonymous user ID for the scan host

ANSWER: A B

Explanation:

BlueXP Classification must be able to mount the export and read its contents, so the export rule must grant the scan host read access. Deep scanning also requires superuser access for the scan host so that the scanner can inspect content regardless of ordinary file ownership and permission restrictions.

Write access is not required for a standard classification scan. Kerberos authentication may be used by an environment, but it is not universally required for BlueXP Classification to scan an NFS export.