

DUMPS ARENA

NetApp Technology Solutions Professional Exam

Netapp NS0-005

Version Demo

Total Demo Questions: 17

Total Premium Questions: 174

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

You are building a new application in Google Cloud that requires multi-protocol file access.

In this scenario, which solution would you use?

- A. Google Cloud Storage
- B. NetApp Astra
- C. NetApp Cloud Volumes Service
- D. Google Cloud Spanner

ANSWER: C

Explanation:

NetApp Cloud Volumes Service is the appropriate solution because it delivers fully managed, cloud-native file storage in Google Cloud with multi-protocol access. A single volume can be configured to serve data through both NFS and SMB, enabling Linux/UNIX and Windows clients to access shared file data while supporting the identity, permissions, and access-control requirements associated with those protocols. This is particularly valuable for applications with heterogeneous compute environments, shared engineering or content repositories, user home directories, and workloads that need enterprise file-service behavior without deploying and operating file servers.

The service is designed to provide persistent, high-performance file storage integrated with Google Cloud. It also includes enterprise storage capabilities such as snapshots, cloning, and data protection features, helping application teams provision and manage shared storage through cloud-oriented workflows. NetApp's Google Cloud documentation identifies Cloud Volumes Service as a managed file service that supports NFS and SMB protocols, while Google Cloud presents it as a managed service for enterprise file workloads. These capabilities directly meet the stated requirement for multi-protocol file access. See the [NetApp Cloud Volumes Service for Google Cloud documentation](#) and the [Google Cloud NetApp Volumes overview](#).

QUESTION NO: 2

Which NetApp product enables the movement of data across different platforms, such as file and object storage?

- A. BlueXP classification
- B. BlueXP copy and sync
- C. Cloud Secure
- D. BlueXP observability

ANSWER: B

Explanation:

BlueXP copy and sync is the NetApp data-movement service designed to transfer and synchronize data among heterogeneous storage locations. It supports moving data between file-based sources and targets, including NFS and SMB shares, and object storage endpoints such as Amazon S3, Azure Blob, Google Cloud Storage, and StorageGRID. The service uses a data broker to securely connect the source and target environments and perform the transfer. It can be used for one-time migrations as well as scheduled or continuous synchronization workflows, making it appropriate when data must be moved across on-premises systems, cloud providers, and object-storage platforms. BlueXP copy and sync is accessed and managed through the BlueXP control plane, where administrators define source and target relationships, configure synchronization settings, and monitor transfer status. This capability directly addresses cross-platform data mobility rather than data discovery, security monitoring, or infrastructure observability. NetApp describes the supported source and target relationships in its [BlueXP copy and sync supported sync relationships documentation](#) and provides operational guidance in the [BlueXP copy and sync documentation](#).

QUESTION NO: 3

Which two protocols are supported with an E-Series system? (Choose two.)

- A. iSCSI
- B. NFS
- C. FC
- D. SMB

ANSWER: A C

Explanation:

NetApp E-Series storage systems are block-storage arrays designed for SAN deployments. **iSCSI** is supported because E-Series controllers can present volumes as block devices across standard Ethernet networks using the iSCSI protocol. This enables host connectivity through IP-based SAN infrastructure while retaining the block-oriented access model used by E-Series volumes. **FC** is also supported because E-Series systems offer Fibre Channel host interfaces for attaching servers through a dedicated Fibre Channel SAN fabric. Hosts access the presented E-Series volumes as block LUNs, and multipath software can provide redundant paths and improved availability. These two connectivity methods are part of the supported host-interface choices for E-Series platforms and are configured through SANtricity storage-management tools. The exact supported interface hardware and speeds depend on the controller model and installed host-interface cards, but iSCSI and Fibre Channel are established E-Series host-connection protocols. NetApp's E-Series documentation describes the supported host-interface configurations and the host-side procedures for connecting through both technologies. See the [NetApp E-Series Linux configuration overview](#) and the [NetApp SANtricity hardware documentation](#).

QUESTION NO: 4

What are two main benefits when configuring Advanced Disk Partitioning (ADP) on an AFF or FAS system? (Choose two.)

- A. Increases deduplication efficiencies.
- B. Reduces the RAID parity overhead.
- C. Decreases disk failure rate.
- D. Optimises storage capacity.

ANSWER: B D

Explanation:

"Reduces the RAID parity overhead" and "Optimises storage capacity" are the two principal benefits of Advanced Disk Partitioning. ADP divides a physical drive into partitions that can be assigned to different aggregates and nodes. In root-data partitioning, each drive provides a small root partition and a larger data partition; root-data-data partitioning can provide one root partition and two data partitions. This design enables a high-availability pair to maintain the required root aggregates without dedicating entire drives solely to root use. Consequently, fewer physical drives are consumed for root and RAID parity purposes, leaving a larger share of the installed drive capacity available for data aggregates. The resulting reduction in the number of drives needed to satisfy aggregate and protection layouts reduces the effective RAID-parity capacity overhead, particularly in smaller AFF configurations where dedicated root drives would otherwise represent a substantial portion of the system. It also improves usable-capacity efficiency while retaining the isolation and aggregate ownership requirements needed by each controller in the HA pair. NetApp describes disk partitioning and its root/data allocation model in its [ONTAP disk and aggregate documentation](#). Additional platform and storage-configuration guidance is available in the [ONTAP disks and aggregates overview](#).

QUESTION NO: 5

Your users require both NFS and SMB file services. You do not want to manage the storage using ONTAP System Manager or the command line interface.

In this scenario, which two products would satisfy these requirements? (Choose two.)

- A. Azure NetApp Files
- B. Cloud Volumes Service
- C. Azure Blob Storage
- D. Amazon S3 Glacier

ANSWER: A B

Explanation:

Azure NetApp Files is a fully managed Azure file-storage service built on NetApp technology. It delivers enterprise NAS capabilities and supports both NFS and SMB volumes, allowing organizations to serve Linux/UNIX and Windows workloads from the same managed service. Administration is performed through Azure management tools, such as the Azure portal, API, and Azure CLI, rather than through ONTAP System Manager or the ONTAP command-line interface. Microsoft documents the available protocols and volume configuration in its [Azure NetApp Files documentation](#).

Cloud Volumes Service is also a fully managed NetApp cloud file service designed to provide shared file storage without requiring customers to operate ONTAP directly. It provides NFS and SMB file-service capabilities for workloads that need cloud-hosted NAS access, while storage provisioning and service administration occur through the cloud service's management experience and APIs. This fits the requirement to avoid direct ONTAP System Manager and CLI administration. NetApp describes the service's managed cloud-storage model and protocol support in the [Cloud Volumes Service documentation](#).

QUESTION NO: 6

Currently, all of your application workloads store their data on a NetApp AFF cluster. You are required to move 20% of your application workloads to the public cloud by the end of the fiscal year. You also want a flexible payment option as you scale your public cloud footprint.

In this scenario, which NetApp storage offering satisfies the requirements?

- A. ONTAP Select deployed on a local vCenter Server instance
- B. Cloud Volumes ONTAP bring-your-own-license (BYOL) subscription
- C. Cloud Volumes ONTAP pay-as-you-go (PAYGO) subscription
- D. NetApp FAS cluster deployed in a secondary location

ANSWER: C

Explanation:

Cloud Volumes ONTAP pay-as-you-go (PAYGO) subscription is the appropriate offering because it extends NetApp ONTAP data management capabilities into supported public-cloud environments while using a consumption-based payment model. This enables the organization to move the required application workloads from its AFF environment to cloud storage without adopting an entirely different storage operating model. Cloud Volumes ONTAP provides familiar ONTAP functionality, including storage efficiency, snapshots, replication, and multiprotocol data access, which helps maintain operational consistency as workloads are placed in the cloud.

The PAYGO model is especially suited to a footprint that will grow over time because charges are based on hourly consumption through the cloud provider's marketplace. Capacity and performance can therefore be adjusted as migration proceeds, rather than requiring a large upfront software commitment. This flexibility supports an incremental migration target and accommodates changing cloud resource needs during and after the fiscal year. NetApp documents Cloud Volumes ONTAP licensing choices, including marketplace-based PAYGO licensing, in its [Cloud Volumes ONTAP licensing](#)

[documentation](#). Additional product capabilities and supported cloud deployment models are described on the [NetApp Cloud Volumes ONTAP product page](#).

QUESTION NO: 7

Which three benefits of using NetApp Cloud Volumes ONTAP for AWS? (Choose three.)

- A. automatic tiering to Amazon S3 for cost reduction
- B. automatic replication to Cloud Volumes Service with SnapMirror
- C. ONTAP thin provisioning for space efficiency and over-provisioning
- D. NetApp Aggregate Encryption (NAE) with Onboard Key Manager
- E. ONTAP data reduction mechanisms with compression and deduplication

ANSWER: A C E

Explanation:

Cloud Volumes ONTAP provides ONTAP storage-efficiency and cloud-tiering capabilities while running in AWS. “automatic tiering to Amazon S3 for cost reduction” is a key benefit because FabricPool tiering can move inactive data blocks from the performance tier to low-cost Amazon S3 object storage according to the configured tiering policy. This lets organizations retain transparent NAS or SAN access through ONTAP while reducing the cost of capacity that does not require high-performance storage.

“ONTAP thin provisioning for space efficiency and over-provisioning” is also a benefit. Thin provisioning presents logical storage capacity to applications while consuming physical capacity only as data is actually written. Combined with monitoring and capacity-management practices, it improves utilization and provides administrators flexibility when allocating cloud storage.

“ONTAP data reduction mechanisms with compression and deduplication” is correct because Cloud Volumes ONTAP supports ONTAP storage-efficiency technologies that reduce the physical capacity required for eligible data. Deduplication eliminates duplicate blocks, while compression reduces the size of stored data blocks. These efficiencies can lower both storage consumption and associated AWS infrastructure costs. See the [NetApp Cloud Volumes ONTAP data-tiering documentation](#) and [ONTAP compression documentation](#).

QUESTION NO: 8

Which storage type provides access using Amazon S3 APIs?

- A. object
- B. NAS
- C. inode
- D. SAN

ANSWER: A

Explanation:

The correct answer is **object**. Object storage organizes data as discrete objects that include the data itself, extensible metadata, and a unique identifier. Rather than mounting a filesystem or presenting block devices to hosts, applications access object data through HTTP-based REST APIs. Amazon Simple Storage Service (Amazon S3) defines the widely adopted S3 API model for this type of access, including operations to create buckets, upload and retrieve objects, list content, and manage object metadata. NetApp StorageGRID is an example of an object-storage platform that provides an S3-compatible API, enabling applications designed for Amazon S3 to store and retrieve data from an on-premises or hybrid-cloud object store. This access model is particularly appropriate for unstructured data, cloud-native applications, backup

repositories, archives, media content, and data lakes. NetApp ONTAP also supports S3 object storage, allowing S3 clients to access data through buckets and S3 API requests. For more details, see the [NetApp StorageGRID S3 REST API documentation](#) and the [NetApp ONTAP S3 configuration documentation](#).

QUESTION NO: 9

A company is running both primary and secondary workloads on a NetApp AFF A250 system. The company wants to free space on their on-premises NetApp ONTAP clusters by migrating inactive data to Amazon Simple Storage Service (Amazon S3).

Which NetApp technology could accomplish this task?

- A. FlexCache
- B. SVM data mobility
- C. BlueXP edge caching
- D. BlueXP tiering

ANSWER: D

Explanation:

BlueXP tiering is the appropriate technology because it extends ONTAP storage to low-cost object storage while retaining the active working set on the AFF A250 system. It uses FabricPool technology in ONTAP to identify inactive data blocks in eligible volumes and tier those blocks to an object-storage tier such as Amazon S3. This releases capacity on the on-premises performance tier without requiring users or applications to change access paths. Data remains logically available through the same ONTAP volumes; if tiered data becomes active again, ONTAP can retrieve it transparently from Amazon S3. BlueXP provides the cloud-based management interface and workflow for discovering ONTAP systems, configuring Amazon S3 as the object store, selecting volumes, and applying an appropriate tiering policy. This makes it well suited to a mixed primary and secondary workload environment where inactive data should be retained but no longer consume expensive local SSD capacity. NetApp documents the supported ONTAP-to-Amazon-S3 tiering architecture and the BlueXP configuration workflow at [Tiering data from on-premises ONTAP clusters to Amazon S3](#). For the underlying ONTAP behavior and tiering policies, see [ONTAP FabricPool documentation](#).

QUESTION NO: 10

An administrator wants Cloud Insights to collect inventory and performance information from on-premises ONTAP systems.

Which two actions are required to collect this infrastructure data? (Choose two.)

- A. Deploy an Acquisition Unit with network access to the ONTAP systems.
- B. Configure an ONTAP data collector to use the Acquisition Unit.
- C. Install an Acquisition Unit directly on each ONTAP controller.
- D. Configure an SNMP trap destination on every ONTAP system.
- E. Deploy a proxy server as the Cloud Insights collection engine.

ANSWER: A B

Explanation:

An **Acquisition Unit** provides the local collection component for Cloud Insights infrastructure monitoring. It must have network connectivity to the ONTAP systems being monitored. The administrator must then configure an appropriate ONTAP data collector and assign it to the Acquisition Unit.

An Acquisition Unit does not need to be installed on every monitored storage system. A proxy server and an SNMP trap destination can be used in particular environments, but neither is a general prerequisite for Cloud Insights infrastructure collection.

QUESTION NO: 11

Which statement is correct about Cloud Volumes ONTAP systems and ONTAP clusters, before you can replicate data between them?

- A. The source volume has been taken offline.
- B. The source and destination volumes are running compatible ONTAP versions.
- C. The volume names on both the source and the destination are the same.
- D. The destination volume has all of the storage efficiency settings disabled.

ANSWER: B

Explanation:

The source and destination volumes are running compatible ONTAP versions. is correct because SnapMirror replication between a Cloud Volumes ONTAP system and an ONTAP cluster depends on ONTAP version interoperability. Before establishing the relationship, the systems must use versions supported for SnapMirror replication with one another. In general, the destination system must run the same ONTAP version as, or a later compatible version than, the source; NetApp publishes the supported interversion relationships in its SnapMirror compatibility guidance. This requirement ensures that the destination can interpret replicated WAFL data, Snapshot copies, and replication metadata generated by the source and can correctly process subsequent updates. Cloud Volumes ONTAP uses the same core ONTAP data-protection technologies as on-premises ONTAP, so version compatibility remains a prerequisite regardless of whether either endpoint is in AWS, Azure, Google Cloud, or a data center. When configuring replication through BlueXP or ONTAP System Manager, administrators should validate the selected source and destination versions and then create the required cluster and SVM peering relationships before initializing the SnapMirror transfer. NetApp documents the version requirements and supported replication configuration workflow in its [BlueXP replication documentation](#) and the [ONTAP supported SnapMirror deployment configurations documentation](#).

QUESTION NO: 12

A manufacturing company wants to build an AI Center of Excellence in their headquarters using data from their remote factories. All data is currently stored on NetApp ONTAP solutions.

Which NetApp technology will enable the company to prepopulate a subset of the data?

- A. Flash Cache
- B. SnapVault
- C. FlexCache
- D. Flash Pool

ANSWER: C

Explanation:

FlexCache is correct because it creates a cache volume that can provide access to data held in an ONTAP origin volume without requiring the entire source dataset to be copied to the headquarters. For an AI Center of Excellence, the organization can create FlexCache volumes near the compute environment and prepopulate them with the particular files or directories required for training, analytics, or validation. This makes the selected working set available locally before users or applications first request it, reducing initial read latency and avoiding the need to transfer all factory data.

ONTAP FlexCache prepopulation is performed with the `volume flexcache prepopulate` command, which supports specifying paths so that only a relevant subset of data is cached. The cache remains associated with its origin and can serve data locally while ONTAP manages cache population and consistency. This approach is well suited to distributed environments in which central teams need fast access to selected remote datasets while keeping the authoritative data on existing ONTAP systems. NetApp documents the prepopulation workflow at [Prepopulate a FlexCache volume](#) and describes the broader architecture at [ONTAP FlexCache documentation](#).

QUESTION NO: 13

What are two Type 1 hypervisors? (Choose two.)

- A. VMware ESXi
- B. Oracle VM VirtualBox
- C. VMware Workstation
- D. Microsoft Hyper-V

ANSWER: A D

Explanation:

VMware ESXi and Microsoft Hyper-V are Type 1, or bare-metal, hypervisors. A Type 1 hypervisor runs directly on the physical server hardware and provides the virtualization layer on which virtual machines and their guest operating systems run. This architecture is designed for server and data-center use, where centralized administration, availability, performance, and hardware resource control are important. VMware describes ESXi as an enterprise-class, bare-metal hypervisor that installs directly on a physical server rather than requiring a general-purpose host operating system. Microsoft Hyper-V likewise uses the Hyper-V hypervisor directly above the hardware; Windows Server or Windows management components operate in the parent partition, while virtual machines operate in child partitions. In NetApp environments, these platforms are commonly used as supported compute hypervisors with shared storage, including SAN and NFS-based datastore designs. Their direct-to-hardware virtualization model is the defining reason they are classified as Type 1 hypervisors. See the [VMware ESXi product information](#) and the [Microsoft Hyper-V technology overview](#).

QUESTION NO: 14

What are three supported deployment solutions for NetApp StorageGRID? (Choose three.)

- A. virtual servers using VMware ESXi
- B. virtual servers using Docker containers
- C. virtual servers using Microsoft Hyper-V
- D. StorageGRID hardware appliances
- E. virtual servers using Amazon EC2

ANSWER: A B D

Explanation:

NetApp StorageGRID supports deployment as virtual grid nodes on VMware ESXi, as purpose-built StorageGRID hardware appliances, and as containerized grid nodes on supported Linux hosts. VMware ESXi is a supported virtualization platform for running StorageGRID virtual machines, enabling grid nodes to be deployed within a vSphere environment. StorageGRID hardware appliances provide an integrated, validated platform in which the appliance supplies the compute, networking, and storage resources required for StorageGRID nodes. StorageGRID's software-based Linux deployment uses Docker containers managed through the StorageGRID host service; therefore, containerized deployment is a supported approach when the underlying host operating system and infrastructure meet StorageGRID requirements. These approaches allow a StorageGRID grid to be built using virtual infrastructure, appliance infrastructure, or supported Linux/container infrastructure,

including combinations of node types where supported. NetApp documents VMware deployment separately from Linux container deployment and appliance installation, reflecting these as the principal supported installation models. See the [StorageGRID VMware installation instructions](#) and the [StorageGRID Red Hat Enterprise Linux installation instructions](#).

QUESTION NO: 15

A company has a large, read-intensive engineering dataset stored on an on-premises ONTAP volume. Remote users need lower-latency access, but the on-premises volume must remain the authoritative copy.

Which two statements describe how FlexCache satisfies these requirements? (Choose two.)

- A. A cache volume can be deployed closer to the remote clients.
- B. The origin volume remains the authoritative data source.
- C. The origin volume is automatically tiered to object storage.
- D. A fully independent writable copy is created at each remote site.
- E. A scheduled block-replication relationship is created for each client.

ANSWER: A B

Explanation:

A FlexCache volume can be placed closer to remote users or applications, allowing requested data to be cached locally and reducing repeated WAN access to the origin. The original ONTAP volume remains the authoritative **origin volume**; FlexCache does not replace it with an independent primary copy.

FabricPool tiers inactive blocks to object storage rather than providing a NAS cache near clients. A FlexClone volume is an independently writable clone, not a distributed cache. SnapMirror creates a replication relationship rather than caching requested data on demand for remote clients.

QUESTION NO: 16

You have multiple on-premises applications writing to ONTAP LUNs and NFS exports. You want to build a disaster recovery solution from your on-premises ONTAP systems to your Microsoft Azure Resource Group while preserving the data efficiency and access properties for block and file.

Which cloud storage destination would you use in Microsoft Azure to accomplish this task?

- A. NetApp Cloud Volumes Service
- B. Azure Blob Destination
- C. Azure NetApp Files
- D. NetApp Cloud Volumes ONTAP

ANSWER: D

Explanation:

NetApp Cloud Volumes ONTAP is the appropriate destination because it runs the ONTAP data-management platform in Microsoft Azure and can be deployed within the customer's Azure environment. It supports replication from on-premises ONTAP using SnapMirror, enabling an established ONTAP disaster-recovery workflow to extend to Azure. SnapMirror transfers changed blocks efficiently and preserves ONTAP volume replication semantics, allowing protected data to retain the characteristics needed for recovery.

Cloud Volumes ONTAP provides both SAN and NAS protocols, including iSCSI for LUN-based workloads and NFS for file workloads. After a disaster or a planned failover, the replicated volumes can be made writable and served to applications

from the Cloud Volumes ONTAP instance. It also retains core ONTAP storage-efficiency capabilities, such as compression, deduplication, and thin provisioning, helping minimize the storage capacity and data-transfer footprint of the DR solution. This combination of ONTAP-to-ONTAP replication, Azure deployment, multiprotocol access, and storage efficiency directly meets the stated block-and-file DR requirement. NetApp documents Cloud Volumes ONTAP replication and disaster-recovery capabilities in the [BlueXP replication documentation](#) and its Azure deployment model in the [Cloud Volumes ONTAP for Azure overview](#).

QUESTION NO: 17

When creating a StorageGRID solution with two sites, how many storage nodes are required?

- A. at least three nodes per site
- B. at least four nodes in total
- C. at least two nodes per site
- D. at least one node per site

ANSWER: A

Explanation:

A two-site StorageGRID solution requires **at least three nodes per site**. This design provides the minimum Storage Node count needed at each site to support resilient, site-aware data placement and availability. StorageGRID uses information lifecycle management (ILM) policies to place replicated or erasure-coded object data across nodes and, when configured, across sites. Having three Storage Nodes at each location gives the grid sufficient independent storage resources to distribute object copies or erasure-coded fragments while retaining operational resilience for planned maintenance and individual node failures.

This minimum also supports a sound foundation for grid expansion, capacity management, and protection objectives that account for an entire-site outage. In a multi-site architecture, StorageGRID administrators should define ILM rules and profiles that deliberately use both sites so that object durability and availability match business requirements. The node count is therefore evaluated per site, rather than only as a combined total across the grid. NetApp documents the minimum of three Storage Nodes for each site in a multi-site StorageGRID deployment. See the [StorageGRID primer](#) and the [StorageGRID erasure coding documentation](#) for architecture and data-protection context.