

Hybrid CloudArchitect Exam

Netapp NS0-604

Version Demo

Total Demo Questions: 8

Total Premium Questions: 80

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Hybrid Cloud Architecture	8
Topic 2, Cloud Data Services	18
Topic 3, Cloud Storage	43
Topic 4, Cloud Operations	11
Total	80

QUESTION NO: 1

A customer has different on-premises workloads with a need for less than 2ms latency.

Which two service levels in NetApp Keystone storage as a service (STaaS) does the customer need? (Choose two.)

- A. Extreme
- B. Standard
- C. Premium
- D. Performance

ANSWER: A C

Explanation:

Extreme and **Premium** are the appropriate NetApp Keystone STaaS performance service levels for on-premises workloads requiring latency below 2 ms. Keystone assigns service levels based on performance characteristics, including the expected latency objective, so the correct selection must meet the stated latency threshold rather than simply provide general-purpose storage performance.

The Extreme service level is intended for the most demanding, latency-sensitive workloads and has a target latency of less than 1 ms. It therefore provides performance comfortably within a less-than-2-ms requirement. The Premium service level has a target latency of less than 2 ms, directly matching the customer's requirement. These tiers support workloads such as high-transaction databases and other business-critical applications whose responsiveness depends on consistently low storage latency.

Latency targets are service-level characteristics and should be validated during sizing together with capacity, IOPS, protocol, and workload-pattern requirements. NetApp documents the service-level performance targets, including the latency objectives for Extreme and Premium, in its Keystone service definitions. See the [NetApp Keystone service levels documentation](#) and the [NetApp Keystone STaaS technical report](#).

QUESTION NO: 2

A company wants a cost-effective storage solution to migrate their VMware environment from on-premises to Azure using Azure VMware Solution. Their current workload requires more storage than compute.

Which datastore storage solution should the company use?

- A. Azure NetApp Files
- B. Azure Files
- C. Amazon FSx for NetApp ONTAP
- D. Cloud Volumes ONTAP in Azure

ANSWER: A

Explanation:

Azure NetApp Files is the appropriate datastore solution because it provides external NFS datastores for Azure VMware Solution clusters. This lets the company scale storage capacity independently of the Azure VMware Solution private-cloud hosts. That separation is particularly valuable for a workload whose storage requirement is larger than its compute requirement: the organization can add the required datastore capacity without purchasing additional host compute and its associated local vSAN storage simply to obtain more space.

Azure NetApp Files is a fully managed Azure service and offers the NFS performance, availability, and VMware integration required for production VMware workloads. Azure VMware Solution can mount Azure NetApp Files volumes as datastores, allowing virtual machines to be migrated or placed on those datastores while retaining VMware operational constructs.

Capacity and service level can be selected to align performance and cost with the workload, making this architecture suitable for storage-intensive migrations.

Microsoft documents Azure NetApp Files datastores as an Azure VMware Solution storage option and describes the configuration requirements and supported NFS connectivity in its [Azure NetApp Files datastore documentation](#). Additional Azure VMware Solution storage architecture guidance is available in the [Azure VMware Solution storage concepts documentation](#).

QUESTION NO: 3

A customer wants to protect an on-premises ONTAP cluster with NetApp BlueXP backup and recovery to Microsoft Azure Blob storage. The customer has deployed a BlueXP Connector in the data center.

Which two connectivity requirements must be included in the design? (Choose two.)

- A. HTTPS connectivity from the Connector to the cluster-management LIF
- B. Outbound HTTPS connectivity from the cluster intercluster LIFs to Azure Blob storage
- C. Inbound HTTPS connectivity from Azure Blob storage to the cluster-management LIF
- D. NFS connectivity from Azure Blob storage to the ONTAP data LIFs
- E. Inbound SSH connectivity from the BlueXP service to every ONTAP node

ANSWER: A B

Explanation:

The **BlueXP Connector must be able to reach the cluster-management LIF over HTTPS** so that BlueXP can discover and manage the on-premises ONTAP cluster and configure the backup service.

The **ONTAP cluster requires outbound HTTPS connectivity from its intercluster LIFs to Azure Blob storage** for backup data transfer. The Connector provides management and orchestration, but it is not the data path for all backup data. Azure Blob storage does not need inbound connectivity to the cluster, and client NFS connectivity is unrelated to the backup workflow.

QUESTION NO: 4

Cyber Monday is quickly approaching as supply chain issues delay a large server shipment to the data center. The VMware environment is storage-heavy and needs to rapidly grow using Azure services.

Which two technologies should the company use to resolve the issue? (Choose two.)

- A. Azure VMware Solution
- B. Amazon FSx for NetApp ONTAP
- C. Google Cloud VMware Engine
- D. Azure NetApp Files

ANSWER: A D

Explanation:

Azure VMware Solution is appropriate because it provides a Microsoft-managed VMware cloud environment running in Azure. It lets an organization extend or migrate existing VMware workloads while continuing to use familiar VMware technologies, including vSphere-based administration and operational processes. This addresses the immediate need for additional compute and VMware capacity without waiting for physical servers to arrive at the data center. Microsoft describes

Azure VMware Solution as a way to run VMware workloads natively on Azure infrastructure and connect them with Azure services.

Azure NetApp Files is the complementary storage technology for a storage-intensive VMware expansion. It provides high-performance, scalable file storage in Azure using NetApp ONTAP technology, and it can provide NFS datastores for Azure VMware Solution. Capacity can be provisioned and scaled in the cloud, allowing the environment to accommodate rapid growth while supplying the low-latency, enterprise-grade shared storage expected by VMware workloads. Together, Azure VMware Solution supplies the VMware software-defined data center capacity, while Azure NetApp Files supplies scalable external datastore capacity for those workloads. See the [Azure VMware Solution documentation](#) and the [Azure NetApp Files datastores for Azure VMware Solution documentation](#).

QUESTION NO: 5

A customer wants to set up disaster recovery in the Central US region for an existing Azure NetApp Files production workload in the East US2 region.

Which feature should the customer use?

- A. cross-zone replication
- B. SnapMirror replication
- C. SyncMirror replication
- D. cross-region replication

ANSWER: D

Explanation:

Azure NetApp Files cross-region replication is the appropriate feature for creating disaster recovery protection between an East US2 production workload and a destination in Central US. This managed replication capability asynchronously transfers changed data from a source Azure NetApp Files volume to a replication destination volume in another Azure region. The destination can be activated when a regional outage or other disaster requires the workload to run from Central US. Replication is configured through a replication relationship that specifies the source and destination volumes and a replication schedule, allowing the customer to align recovery-point objectives with business requirements. Azure NetApp Files cross-region replication is built on NetApp SnapMirror technology, but “cross-region replication” is the Azure NetApp Files feature name and is the configuration workflow used for this scenario. During a failover, the destination replication relationship is broken so that the destination volume becomes writable and clients can access the recovered workload. After the source region is available again, the relationship can be reestablished and resynchronized for failback. For planning requirements, both regions must support Azure NetApp Files and the destination capacity and networking must be prepared to host the recovered application. See [Azure NetApp Files cross-region replication overview](#) and [cross-region replication requirements and considerations](#).

QUESTION NO: 6

A customer wants to add personal data identifiers from an Oracle database to their NetApp BlueXP classification scans.

Which mechanism should the customer use?

- A. RegEx
- B. custom categories
- C. Data Fusion
- D. custom keywords

ANSWER: C

Explanation:

Data Fusion is the appropriate mechanism because it enables BlueXP classification to incorporate an organization's own list of personal-data identifiers into classification results. A customer can export relevant identifiers maintained in an Oracle database, such as employee IDs, customer numbers, account identifiers, or other unique values, and provide that data to BlueXP classification through Data Fusion. BlueXP classification then uses the supplied identifiers as additional context while scanning supported data sources, helping identify files and records that contain those organization-specific values.

This is useful when the identifiers that matter to the business are not limited to standard, built-in personal-data patterns. Data Fusion extends the classification service with customer-provided reference data, allowing the organization to locate sensitive information based on values that are meaningful in its own environment. The Oracle database is therefore the source of the identifier list, while Data Fusion is the BlueXP classification capability that brings those identifiers into the scan process. NetApp documents Data Fusion as the feature for adding personal data identifiers to scans in the [BlueXP classification custom classification documentation](#). Additional product documentation is available in the [BlueXP classification documentation library](#).

QUESTION NO: 7

A customer deploys NetApp Cloud Volumes ONTAP in AWS and tiers inactive data to an Amazon S3 bucket. The customer does not allow Internet or NAT gateway access from the Cloud Volumes ONTAP subnets. Access through the endpoint must be restricted to the tiering bucket.

Which two configurations should the customer include? (Choose two.)

- A. Associate an Amazon S3 gateway endpoint with the Cloud Volumes ONTAP subnet route tables.
- B. Create an endpoint policy that permits the required S3 actions for the tiering bucket.
- C. Add an inbound security-group rule from the S3 public IP ranges.
- D. Attach an Internet gateway to each Cloud Volumes ONTAP subnet.
- E. Configure an Elastic IP address for each data LIF.

ANSWER: A B

Explanation:

An **Amazon S3 gateway endpoint associated with the Cloud Volumes ONTAP subnet route tables** provides private connectivity from the VPC to the regional S3 service. The route-table association directs applicable S3 traffic through the gateway endpoint instead of a NAT gateway or Internet gateway.

An **endpoint policy that permits the required S3 actions for the tiering bucket** restricts endpoint use to the intended bucket and operations. A security group does not control traffic through an S3 gateway endpoint, and an Internet gateway would violate the requirement to avoid public Internet access.

QUESTION NO: 8

A customer is implementing NetApp StorageGRID with an Information Lifecycle Management (ILM) policy. Which key benefit should the customer expect from using ILM policies in this solution?

- A. improved data security
- B. automated data optimization
- C. real-time data analytics capabilities
- D. simplified data access controls

ANSWER: B

Explanation:

Automated data optimization is the key benefit of StorageGRID Information Lifecycle Management policies. StorageGRID uses ILM rules and policies to automatically evaluate objects and apply the defined instructions for where, how, and for how long object data is stored. Administrators can define placement instructions based on object metadata, such as bucket, tenant, object name, size, or creation time. The policy can then place copies on appropriate storage pools or cloud storage pools and maintain the required level of protection through replication or erasure coding. As objects age or meet specified criteria, StorageGRID automatically performs the policy-directed placement and retention actions without requiring an administrator to manually move individual objects. This enables the organization to align storage consumption, data durability, geographic placement, and retention requirements with the value of data throughout its lifecycle. Proper ILM design therefore provides policy-driven, automated optimization of object storage resources while maintaining the intended data-protection posture. NetApp describes ILM as the framework that manages object data over time through rules and policies; see the [StorageGRID ILM documentation](#) and the [ILM rules overview](#).