

NetApp StorageGRID Administrator Exam

Netapp NS0-077

Version Demo

Total Demo Questions: 7

Total Premium Questions: 73

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, StorageGRID Architecture	12
Topic 2, StorageGRID Installation	4
Topic 3, StorageGRID Configuration	32
Topic 4, StorageGRID Administration	18
Topic 5, StorageGRID Troubleshooting	7
Total	73

QUESTION NO: 1

A StorageGRID load balancer endpoint uses the DNS name s3.example.com. Clients use virtual-hosted-style requests for buckets such as reports.s3.example.com and media.s3.example.com. The endpoint certificate contains only s3.example.com, and clients report TLS hostname validation errors.

Which certificate configuration resolves the issue?

- A. A certificate containing *.s3.example.com.
- B. A certificate containing only s3.example.com.
- C. A certificate containing the Grid Manager URL.
- D. A certificate containing the IP addresses of all Storage Nodes.

ANSWER: A

Explanation:

A certificate containing *.**s3.example.com** resolves the hostname mismatch for virtual-hosted-style bucket requests. The wildcard name matches a single bucket-name label before the endpoint domain, such as reports.s3.example.com.

A certificate containing only s3.example.com is valid for path-style access to that endpoint name, but it does not match bucket-specific host names. Changing the bucket consistency level or creating additional tenant accounts has no effect on TLS hostname validation.

QUESTION NO: 2

To comply with an industry regulation, a bank is establishing a disaster recovery solution across two grids. The tenant federation was successfully configured. You are required to unidirectionally replicate a critical bucket from Grid-1 to Grid-2 and ensure that all of the content is identical on each grid, to reestablish operations in the event that Grid-1 fails.

After creating a bucket with versioning enabled on each grid, what are the required steps to complete the task?

- A. Enable cross-grid replication (CGR) on each grid, select the grid federation connection, and select Yes for Replicate delete markers.
- B. Enable cross-grid replication (CGR) on Grid-2, select the grid federation connection, and select Yes for Replicate delete markers.
- C. Enable cross-grid replication (CGR) on Grid-1, select the grid federation connection, and select Yes for Replicate delete markers.
- D. Enable cross-grid replication (CGR) on Grid-1, select the grid federation connection, and select Yes for Replicate metadata.

ANSWER: C

Explanation:

Enable cross-grid replication (CGR) on Grid-1, select the grid federation connection, and select Yes for Replicate delete markers. Cross-grid replication is configured on the source bucket, and it automatically replicates eligible object operations from that bucket to its corresponding bucket on the other grid. Since the required disaster-recovery flow is from Grid-1 to Grid-2, Grid-1 must be the source bucket where CGR is enabled. The grid federation connection selected during configuration identifies the paired grid and matching destination bucket that receive the replicated versions.

Both buckets must have versioning enabled before replication is configured because StorageGRID uses S3 versioning semantics to replicate object versions and deletion activity. Selecting *Yes for Replicate delete markers* ensures that a delete marker created in the Grid-1 source bucket is replicated to Grid-2. This keeps the visible object state aligned between the two buckets when clients perform ordinary S3 deletes without specifying a version ID. Consequently, Grid-2 maintains a synchronized recovery copy suitable for assuming operations after a Grid-1 failure.

NetApp's cross-grid replication documentation describes configuring replication from the source bucket and the handling of replicated delete markers. See [Manage cross-grid replication](#) and [Create an S3 bucket](#).

QUESTION NO: 3

Consider a three-site StorageGRID deployment with three Storage Nodes per site. The consistency level for a bucket is set to strong-global.

Which two statements explain the read and write behavior on this bucket? (Choose two.)

- A. READ returns the record after a quorum of replica nodes on each site have responded.
- B. WRITE must be written to a quorum of replica nodes in each site.
- C. WRITE must be written to a quorum of replica nodes in the local site.
- D. READ returns the record after a quorum of replica nodes in the local site have responded.

ANSWER: B D

Explanation:

With the legacy StorageGRID *strong-global* consistency model assumed by this question, a successful write requires metadata-replica quorum acknowledgement at every site. Each site has three Storage Nodes, so the replica quorum at an individual site is two nodes. Therefore, the write is acknowledged only after two metadata replicas at each of the three sites have recorded the update. This behavior ensures that the object metadata update is established across the grid before the client receives a successful response.

A read under strong-global requires a quorum response from metadata replicas at the client's local site. In this deployment, that means two local metadata replicas must respond before StorageGRID returns the record. Combining grid-wide write acknowledgement with a local read quorum provides read-after-write consistency across sites while allowing read requests to be served through the local site's metadata replicas.

NetApp documents strong-global consistency as the setting intended to provide read-after-write consistency for requests across all sites. The detailed legacy behavior described here is applicable to the StorageGRID release semantics represented by this certification question; newer releases can use updated quorum strong-global behavior. See the NetApp documentation on [S3 consistency controls](#) and the [StorageGRID architecture overview](#).

QUESTION NO: 4

An administrator has created a new ILM policy for a tenant workload. Before activating the policy, the administrator must verify that objects with a particular bucket name, object-key prefix, and user metadata value will match the intended rule.

Which method should be used?

- A. Use the ILM policy simulation tool with representative object attributes.
- B. Review ORLM audit messages for an object that will be ingested after activation.
- C. Perform a metadata lookup to confirm the object location.
- D. Create a traffic classification policy for the tenant workload.

ANSWER: A

Explanation:

Use the ILM policy simulation tool with representative object attributes. The simulation evaluates the proposed policy in rule order and shows which rule would match an object with the specified bucket, key, metadata, size, or other applicable filter criteria. It allows the administrator to validate the expected rule before the policy is activated.

An ORLM audit message can identify the rule that was applied to an object after ingest and ILM evaluation has occurred, but it cannot validate a proposed policy in advance. Reviewing object location also shows an outcome rather than the policy rule that would match a prospective object.

QUESTION NO: 5

A tenant is permitted to use StorageGRID platform services. The tenant must asynchronously replicate objects from a source bucket to an external S3-compatible object store immediately after objects are created.

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

- A. Create a CloudMirror endpoint for the external object store.
- B. Add a replication configuration to the source bucket.
- C. Create an event-notification configuration on the source bucket.
- D. Add the external object store as a Cloud Storage Pool in an ILM rule.

ANSWER: A B

Explanation:

Create a CloudMirror endpoint for the external object store and **add a replication configuration to the source bucket** are required. The endpoint defines the destination connection details, while the bucket replication configuration identifies the objects that StorageGRID should mirror to that destination.

Enabling platform services is a grid-level tenant permission that is already present in this scenario. An event-notification configuration can notify an application of object activity, but it does not replicate objects. An ILM rule controls StorageGRID object placement and is not a CloudMirror replication configuration.

QUESTION NO: 6

You are installing a NetApp SG5860 appliance with the Grid Network and Client Network enabled.

Which two configurations are supported for the Grid and Client networks using the four interfaces? (Choose two.)

- A. A two-port active-backup bond without VLAN tagging.
- B. A four-port active-backup bond with VLAN tagging.
- C. A four-port LACP bond with VLAN tagging.
- D. A two-port LACP bond without VLAN tagging.

ANSWER: A C D

Explanation:

The SG5860 uses the SG5800 appliance networking model, which provides four high-speed interfaces and supports both Fixed and Aggregate port-bond modes. A two-port active-backup bond without VLAN tagging is supported in Fixed mode. Fixed mode assigns one pair of appliance ports to the Grid Network and the other pair to the Client Network. Each pair can use Active-Backup bonding, and VLAN tags are optional when the networks use their respective physical port pairs.

A two-port LACP bond without VLAN tagging is also supported in Fixed mode. The Grid Network pair and Client Network pair can each be configured as an LACP bond, provided that the corresponding switch ports are configured for LACP. As with Active-Backup in Fixed mode, VLAN tagging is optional because the two StorageGRID networks are carried over separate port pairs.

A four-port LACP bond with VLAN tagging is supported through Aggregate port-bond mode. Aggregate mode combines all four interfaces into one LACP bond. Because both the Grid Network and Client Network traverse that shared bond, each

network must have a distinct VLAN tag to keep the traffic separated. NetApp documents these Fixed-mode and Aggregate-mode configurations for SG5800-series appliances. Therefore, the question's instruction to choose only two is ambiguous: all three listed supported configurations should be selected.

See [NetApp StorageGRID Appliances: Configure network links for SG5800](#) and [NetApp StorageGRID Appliances: Port bond modes](#).

QUESTION NO: 7

A company is scaling its StorageGRID environment to support a massive increase in object count, aiming for trillions of objects over the next few years. They need an appliance model that can handle this significant increase in scale and performance demands.

Which NetApp StorageGRID appliance model takes advantage of the increased scaling limits?

- A. SG5812
- B. SG1100
- C. SG5860
- D. SG6160

ANSWER: D

Explanation:

SG6160 is the appropriate appliance for a StorageGRID deployment that requires substantial scale and high performance as object counts grow. It is the high-capacity model in the SG6100 appliance family and combines an SG6100-CN compute controller with a 60-drive storage shelf. Its architecture is intended for demanding object-storage workloads, including large data lakes and workloads with many small or transactional objects, where both metadata-related processing and storage capacity are important.

A key scaling advantage of the SG6160 is its support for storage expansion shelves. One appliance can use its base 60-drive shelf plus up to two additional 60-drive expansion shelves, enabling up to 180 drives per appliance. This permits significant capacity growth without introducing a separate appliance for each additional shelf set. The SG6100-CN also supplies the processing and memory resources needed for StorageGRID node services in large deployments.

StorageGRID is designed to scale horizontally by adding nodes and sites, while the SG6160 provides a strong capacity-and-performance building block within that grid. NetApp documents the SG6160 expansion configuration and intended large-scale use cases in its [SG6100 appliance documentation](#) and describes StorageGRID scalability concepts in the [StorageGRID primer](#).