
Exam : NS0-077

**Title : NetApp StorageGRID
Administrator Exam**

Version : DEMO

1. You are setting up a StorageGRID environment that must maintain data for regulatory compliance. Which two items are required? (Choose two.)

- A. default SnapLock policy
- B. Compliance Clock
- C. S3 Object Lock
- D. compliant default ILM rule

Answer: C D

Explanation:

For regulatory retention in StorageGRID, the required mechanisms are S3 Object Lock and a compliant default ILM rule. StorageGRID implements write-once-read-many protection through S3 Object Lock. NetApp states that the global S3 Object Lock setting must be enabled before tenants can create buckets with S3 Object Lock enabled. Once enabled globally, the setting cannot be disabled.

There is also a mandatory ILM dependency. NetApp explicitly requires the default rule in every active ILM policy to be compliant before global S3 Object Lock can be enabled. A compliant default rule must maintain at least two replicated copies or one erasure-coded copy on Storage Nodes, must begin protection at day 0 using Ingest time, cannot place the protected copies in a Cloud Storage Pool, and must include a placement extending forever.

Option A is incorrect because SnapLock is an ONTAP WORM technology and is not the StorageGRID compliance mechanism.

Option B is also incorrect: StorageGRID regulatory retention does not require an administrator-configured "Compliance Clock." Object retention is enforced through S3 Object Lock retain-until dates, retention modes, legal holds, and compliant ILM placement.

Reference topics: S3 Object Lock # Regulatory Compliance # Compliant ILM Rules # Default ILM Rule # Retention and Legal Hold

2. Which encryption method is available for all deployment methods?

- A. Node Encryption
- B. Drive Security
- C. Stored Object Encryption
- D. Drive Encryption

Answer: C

Explanation:

Stored Object Encryption is the StorageGRID encryption method that is available regardless of how the grid nodes are deployed. It is configured centrally in Grid Manager rather than depending on appliance-specific hardware capabilities. When enabled, StorageGRID encrypts newly ingested S3 object data that has not already been encrypted at the bucket or object level. StorageGRID supports AES-128 or AES-256 for this function.

The other options are deployment-specific. Node Encryption is associated with StorageGRID appliance nodes and a configured external Key Management Server (KMS); it must be enabled for an eligible appliance during installation. Drive Security depends on storage appliances equipped with Full Disk Encryption or self-encrypting drives and is managed through SANtricity System Manager. Drive Encryption likewise requires StorageGRID appliances containing supported self-encrypting drives and is configured through the StorageGRID Appliance Installer. These hardware dependencies mean none of those three methods is universally available across virtual, containerized, cloud, and appliance

deployment models.

Therefore, C. Stored Object Encryption is the correct answer because it operates at the StorageGRID object-ingest layer rather than relying on a particular node or disk hardware implementation.

Reference topics: Grid Manager # Security Settings # Network and Objects # Stored Object Encryption; StorageGRID Encryption Methods

3.Which StorageGRID feature enables you to have multiple active information lifecycle management (ILM) policies?

- A. bucket tags
- B. ILM policy tags
- C. multitenancy
- D. Identity and Access Management (IAM) policy

Answer: B

Explanation:

The feature that enables multiple active ILM policies is ILM policy tags. NetApp StorageGRID allows administrators to create multiple ILM policies and associate each policy with a specific policy tag. Tenants can then assign one of those ILM policy tags to an S3 bucket, causing objects in that bucket to be managed by the ILM policy associated with the selected tag. This enables different buckets to receive different data-protection, placement, retention, or performance treatments while those policies remain active simultaneously.

StorageGRID also maintains a Default ILM policy tag. Every grid must have an active policy associated with the Default tag, and that policy governs objects in buckets that do not have another ILM policy tag assigned. NetApp supports up to 10 policy tags per grid.

Option A is too general. Although StorageGRID implements the assignment on an S3 bucket, the enabling construct is specifically the ILM policy tag.

Option C, multitenancy, separates tenant resources but does not itself permit multiple active ILM policies.

Option D, IAM policy, controls S3 access permissions and authorization rather than ILM placement or retention behavior.

Therefore, B is the correct answer.

Reference topics: ILM # Policy Tags # Multiple Active Policies # Default Policy Tag # Per-Bucket ILM Assignment

4.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:

For unidirectional cross-grid replication, CGR must be enabled on the bucket in the grid that acts as the source. NetApp states that when CGR is enabled on only one grid, objects added to that source bucket are replicated automatically to the corresponding destination bucket on the other grid; objects written to the destination are not replicated back. Because the required direction is Grid-1 # Grid-2, CGR must therefore be enabled on Grid-1.

The administrator then selects the appropriate grid federation connection and sets Replicate delete markers = Yes. When an S3 client deletes an object without specifying a version ID, StorageGRID creates a delete marker. With this option enabled, that marker is replicated to Grid-2 so the object appears deleted on both grids, helping maintain matching bucket state. NetApp documents this as the default behavior.

Option A would configure bidirectional replication, which violates the unidirectional requirement.

Option B reverses the required replication direction.

Option D is incorrect because the relevant CGR configuration control is Replicate delete markers, not "Replicate metadata."

Reference topics: Tenant Manager # Grid Federation # Cross-Grid Replication # Source and Destination Buckets # Versioning # Replicate Delete Markers

5.A StorageGRID ILM policy has moved objects belonging to a specific bucket to a Cloud Storage Pool. The objects have transitioned automatically to S3 Glacier Deep Archive storage class to minimize costs. A user now requires immediate access to one of these archived objects.

How can the user achieve this goal?

A. The user must issue an S3 GetObject request, and StorageGRID will automatically retrieve the object.

B. Since the object is no longer on StorageGRID, the user must go directly to the Amazon S3 console to restore and retrieve the object.

C. Retrieval is not possible, because StorageGRID does not support retrieval from S3 Glacier Deep Archive.

D. The user must issue an S3 RestoreObject request from Glacier Deep Archive, followed by a GetObject request.

Answer: D

Explanation:

When an object stored in a Cloud Storage Pool has transitioned to Amazon S3 Glacier Deep Archive, it is in a non-retrievable state. StorageGRID therefore cannot satisfy a normal GetObject request until a retrievable copy has first been restored. NetApp specifies that the client application must issue an S3 RestoreObject request to restore the object temporarily within the Cloud Storage Pool. After the restoration completes, the client can issue GetObject through StorageGRID to retrieve the data.

For a restored object, RestoreObject specifies how long the retrievable copy remains available.

StorageGRID supports restoration from Glacier Deep Archive, but Deep Archive does not support the Expedited retrieval tier; only Standard or Bulk retrieval can be used. Thus, "immediate" access means initiating the restore process immediately, not instantaneous retrieval from the Deep Archive tier.

Option A is incorrect because GetObject does not automatically restore a non-retrievable Glacier object.

Option B is incorrect because restoration is performed through StorageGRID's S3 API; direct Amazon-console intervention is unnecessary.

Option C is incorrect because StorageGRID explicitly supports RestoreObject for Cloud Storage Pool objects transitioned to Glacier Deep Archive.

Reference topics: ILM # Cloud Storage Pools # S3 Glacier Deep Archive # RestoreObject # GetObject # Retrieval Tiers