



Vendor: Network Appliance

Exam Code: NS0-077

Exam Name: NetApp StorageGRID Administrator Exam

Version: DEMO

QUESTION 1

Your organization is deploying StorageGRID across multiple data centers in different geographic regions. You need to ensure that object data is distributed across these regions for disaster recovery, while also maintaining local copies for performance. Which combination of StorageGRID components would you configure to achieve this requirement?

- A. Create multiple storage pools associated with different regions, then define ILM rules that place objects in specific pools based on ingest location
- B. Configure a single storage pool and use VLAN interfaces to automatically route traffic to the nearest data center
- C. Deploy Load Balancer endpoints across all regions and configure traffic classification rules to distribute objects
- D. Use only high-availability groups and configure them to replicate data across regions without creating separate storage pools

Answer: A

Explanation:

The correct approach is to create multiple storage pools and define ILM rules that reference these pools by region. Storage pools are the fundamental mechanism for controlling where object data is stored in StorageGRID, and they can be associated with specific locations. ILM rules then place objects into these pools based on criteria such as ingest location, object metadata, or time-based policies. This directly addresses the disaster recovery and performance requirements.

While VLAN interfaces, load balancer endpoints, and high-availability groups are all valid StorageGRID components, they serve different purposes. VLANs are for network segmentation, load balancer endpoints manage client connections, and HA groups provide failover for gateway nodes--none of these directly control object placement across regions the way storage pools and ILM rules do.

QUESTION 2

A StorageGRID site is moving to a new data center with a different network subnet. The environment will remain online with information lifecycle management (ILM) policies that ensure data is replicated across sites. What is used to configure the network settings for the new data center?

- A. StorageGRID Appliance Installer
- B. change-ip tool
- C. ifconfig command on each node
- D. Tenant Manager

Answer: B

Explanation:

The correct tool for modifying network addressing on an existing StorageGRID deployment is the change-ip tool. NetApp documents that the Change IP tool can modify the network configuration of one or more grid nodes, including the IP address, subnet mask/prefix, gateway, and MTU. It can also update Grid Network configuration and add, modify, or remove Admin and Client Network addressing.

This is the appropriate mechanism when moving an operational StorageGRID site to a new data center where the network subnet changes. The tool provides a controlled workflow to select affected nodes, edit addressing, review the proposed configuration, validate the changes, and then apply them to the grid. StorageGRID 12.0 documentation shows the workflow explicitly through the change-ip command and its menu options for editing IP/mask, gateway, MTU, and Grid Network subnet information.

QUESTION 3

Which of the following scenarios is a use case and benefit for leveraging the integrated S3 caching layer in NetApp StorageGRID?

- A. accelerating performance and reducing latency for read-heavy, frequently accessed S3 data, especially in geographically distributed deployments
- B. reducing long-term data archival costs by automatically moving infrequently accessed data to public cloud cold storage tiers
- C. providing highly available, on-premises object storage for backup and disaster recovery operations that prioritize data durability and integrity over access speed
- D. enabling multiprotocol (for example, S3, SMB, NFS) access to the same dataset on a single volume within a NetApp ONTAP cluster for seamless collaboration

Answer: A

Explanation:

StorageGRID's integrated S3 caching capability is designed to improve performance for workloads that repeatedly read a working subset of object data. NetApp states that load balancer caching can significantly improve performance when a workload operates on a subset of data and accesses objects multiple times. The cache is implemented on Gateway Nodes and allows frequently requested S3 objects to be served locally rather than repeatedly retrieving the data from upstream Storage Nodes. This reduces retrieval latency and backend traffic, which is particularly valuable when applications or compute resources are geographically separated from the primary object-storage location.

NetApp identifies data lakes, high-performance computing, AI/ML training, content distribution networks, media asset management, and video production as suitable caching workloads. StorageGRID 12.0 specifically highlights the S3 caching layer as a performance enhancement for AI/ML training.

QUESTION 4

You are configuring a two-site StorageGRID solution to be the destination for mission-critical application data. You want the HA group to fail over from one site to the other in the event of a site failure. How is the HA group configured?

- A. Enter a virtual IP address in a subnet that spans all sites in the virtual IP address list.
- B. Enter a unique IP address for each site in the virtual IP address list.
- C. Enter a fully qualified domain name (FQDN) for the virtual IP address.
- D. Enter a gateway IP address for the virtual IP address.

Answer: A

Explanation:

A StorageGRID High Availability (HA) group provides active-backup access by assigning one or more virtual IP (VIP) addresses to the currently active interface. For site-level failover, the HA group includes suitable Admin Node or Gateway Node interfaces from both sites, and those interfaces must participate in a network design that allows the same VIP to move between sites. NetApp's StorageGRID site-failover guidance specifically describes including interfaces from Site 1 and Site 2 and assigning a VIP within the same subnet. If the preferred site's interface becomes unavailable, StorageGRID transfers the VIP to the next available interface at the other site.

Accordingly, the VIP subnet must be available across the participating sites; the same client-facing VIP remains usable after failover. StorageGRID requires the VIP and optional gateway to belong to the configured VIP subnet.

QUESTION 5

You have a StorageGRID system that spans multiple sites. You are asked to create an S3 bucket for extremely important objects that require all sites to receive metadata immediately after any new write or change. When creating this S3 bucket, which consistency level meets the requirement?

- A. strong-global
- B. available
- C. strong-site
- D. read-after-new-write

Answer: A

Explanation:

The appropriate consistency level is strong-global because the requirement explicitly applies to clients and object metadata across multiple StorageGRID sites. NetApp defines strong-global as providing read-after-write consistency for all client requests across all sites. Therefore, after an object is created or its metadata is changed, subsequent requests from any site receive a consistent view of that change.

strong-site is insufficient because its consistency guarantee applies only to client requests within a single site. A client accessing the bucket from another StorageGRID site could therefore fall outside the required site-wide consistency guarantee. read-after-new-write, which is the default and generally recommended setting, guarantees consistency for newly created objects but provides eventual consistency for object updates. Because the scenario requires immediate consistency following both new writes and changes, it does not meet the stated requirement. available provides eventual consistency for both new objects and updates and prioritizes availability rather than cross-site consistency.

Thus, among the available options, strong-global is specifically designed for workloads requiring globally consistent object and metadata visibility across a multisite StorageGRID deployment.

QUESTION 6

Your customer is running a multisite grid and is ingesting log data from a variety of sources. The logs are processed by a Security Information and Event Management (SIEM) solution and stored in the grid for reading purposes only. A log file must be readable from all available nodes in the grid. Ingest into the grid must be possible even when a site is down. Due to the size and importance, the log files are stored in an EC7+5 format across three sites. Which bucket consistency level and ingest policy must be selected?

- A. read-after-new-write and strict
- B. strong-global and strict
- C. all and balanced
- D. read-after-new-write and balanced

Answer: D

Explanation:

The correct combination is read-after-new-write consistency with Balanced ingest. The workload contains log objects that are written once and subsequently read, so read-after-new-write provides the required behavior: StorageGRID guarantees read-after-write consistency for new objects, while maintaining high availability. NetApp recommends this consistency setting for most workloads. Because these SIEM log objects are not expected to be updated after creation, the eventual-consistency behavior for updates is not material to this use case.

The ingest requirement is decisive. The EC7+5 placement spans three sites, but clients must remain able to ingest when one site is unavailable. With Strict ingest, StorageGRID must satisfy the ILM placement instructions synchronously before acknowledging the write. If the required EC placement cannot be created because a site is down, the request fails. In contrast, Balanced attempts the specified ILM placement immediately but, when an outage prevents that placement, StorageGRID falls back to dual-commit protection and acknowledges the object. ILM subsequently creates the required EC7+5 placement when resources become available again.

all is inappropriate because it requires every metadata node to receive metadata immediately, reducing availability during failures. Therefore, D satisfies both immediate visibility of newly written logs and continued ingest during a site outage.

QUESTION 7

A primary interface in an HA group has failed, and the virtual IP address (VIP) has moved to a backup interface. What happens when the primary interface is restored?

- A. The Load Balancer service is restarted, and VIP returns to the primary interface.
- B. The VIP remains with the backup interface.
- C. The Load Balancer service is restarted, and VIP remains on the backup interface.
- D. The VIP automatically returns to the primary interface.

Answer: D

Explanation:

StorageGRID HA groups use an explicit interface priority order. The first interface in that order is the Primary interface and normally hosts all virtual IP addresses assigned to the HA group. If the Primary interface becomes unavailable, StorageGRID automatically moves the VIP addresses to the highest-priority backup interface that remains operational.

Importantly, this behavior includes automatic failback. NetApp states that when the failure is resolved and a higher-priority interface becomes available again, the VIP addresses are automatically moved to the highest-priority available interface. Therefore, once the original Primary interface is restored and considered healthy, StorageGRID automatically transfers the VIP back to it. No administrator intervention or Load Balancer service restart is required.

QUESTION 8

An enterprise is expanding a single-site StorageGRID deployment and is adding new nodes on a new subnet. What is a requirement for configuring the Grid Network?

- A. Each node must have a static route that is configured manually for the new subnet.
- B. A Gateway Node must be deployed to manage traffic between the subnets.
- C. The new subnet must be added to the Grid Network subnet list.
- D. All nodes on the Grid Network must be on the same subnet.

Answer: C

Explanation:

When an existing StorageGRID deployment is expanded with nodes whose Grid Network IP addresses are on a subnet that has not previously been used, that subnet must be added to the Grid Network subnet list before the expansion is started. StorageGRID maintains this list so that every grid node can correctly route internal Grid Network traffic to the other participating subnets. NetApp StorageGRID 12.0 documentation explicitly states that if any new node uses a new Grid Network subnet, the administrator must add that subnet first; otherwise, the expansion must be cancelled, the subnet added, and the expansion restarted.

The configuration is performed in Grid Manager Maintenance Network Grid Network, where the administrator selects Add another subnet and enters the network in CIDR notation. StorageGRID then distributes the routing configuration throughout the grid.

QUESTION 9

You have a StorageGRID environment consisting of two sites deployed as virtual machines. Site A contains the primary Admin Node, one Gateway Node, and four Storage Nodes. Site B contains a nonprimary Admin Node, one Gateway Node, and four Storage Nodes. All client operations use load balancer endpoints through the Gateway Node. A host in Site A running both the primary Admin Node and Gateway Node has crashed. What recovery process should you use to resume client operations the quickest?

- A. Deploy a new primary Admin Node, and expand as a new primary Admin Node.
- B. Deploy a new Gateway Node, and recover the failed Gateway Node from the nonprimary Admin Node.
- C. Deploy a new primary Admin Node, and recover the failed primary Admin Node.
- D. Deploy a new Gateway Node, and expand as a new Gateway Node.

Answer: C

Explanation:

When a physical or virtual host containing multiple StorageGRID nodes fails, NetApp defines a specific recovery-order requirement when one of those nodes is the primary Admin Node. StorageGRID 12.0 documentation states that if the failed server hosts the primary Admin Node, the primary Admin Node must be recovered first. This prevents subsequent node-recovery operations from stopping while attempting to communicate with services associated with the primary Admin Node.

Therefore, the administrator should deploy a replacement VM for the primary Admin Node and perform the documented primary Admin Node recovery. After the primary Admin Node is operational, recovery of the failed Gateway Node can proceed, restoring the intended client load-balancing path.

QUESTION 10

Your organization is integrating an on-premises ONTAP cluster with StorageGRID using FabricPool. You need to ensure that frequently accessed data remains on the ONTAP system while less-frequently accessed data is tiered to StorageGRID, and you want to retrieve the tiered data when needed. Which StorageGRID and FabricPool features must be properly configured to support this use case? (Select two)

- A. Configure load balancer endpoints on StorageGRID with SSL/TLS certificates that match the FabricPool endpoint DNS names
- B. Enable the Platform Services feature on the StorageGRID tenant account so that ONTAP can invoke notifications
- C. Create an S3 bucket with appropriate access permissions and ensure the FabricPool service account has sufficient credentials to read and write objects
- D. Configure ILM policies to automatically migrate data back to ONTAP after a specified retention period
- E. Set up VLAN interfaces on StorageGRID to isolate FabricPool traffic from other tenant traffic for security

Answer: AC

Explanation:

Two configurations are critical for FabricPool integration: (1) Load balancer endpoints with proper SSL/TLS certificates that match the DNS names ONTAP uses to connect to StorageGRID, ensuring secure and predictable connectivity; and (2) An S3 bucket with appropriate permissions and valid credentials for the FabricPool service account to perform read and write operations.

Platform Services is used for notifications and transformations, not for tiering operations. ILM policies control StorageGRID object management, not repatriation back to ONTAP--FabricPool handles tiering decisions based on access patterns, and ONTAP retrieves data on demand. VLAN interfaces, while useful for network segmentation, are not required for basic FabricPool functionality and do not directly enable tiering.

QUESTION 11

A customer has installed a hotfix on their StorageGRID system. After the update, an "ILM placement unachievable" alert is reported by the system. What does the alert mean?

- A. POST operations failed, because the information lifecycle management (ILM) rule could not be fulfilled.
- B. GET operations failed, because no copy of the object was available.
- C. PUT operations failed, because the information lifecycle management (ILM) rule could not be fulfilled.
- D. PUT operations could not be stored as defined in the information lifecycle management (ILM) rule.

Answer: D

Explanation:

The ILM placement unachievable alert means that StorageGRID is unable to satisfy one or more placement instructions defined by an active ILM rule for certain objects. NetApp defines this alert as indicating that an ILM placement instruction cannot be achieved for certain objects, commonly because a required Storage Node or storage location is temporarily unavailable or because the ILM rule specifies a placement that the current topology cannot satisfy.

Therefore, D is the most accurate answer. The alert does not inherently mean that every S3 PUT operation failed. Whether a client PUT fails depends on the ILM rule's ingest behavior. With Strict ingest, StorageGRID must create all placements specified by the ILM rule immediately, so a placement failure can cause the client PUT to fail. With Balanced ingest, StorageGRID can temporarily use alternate protection and complete the required ILM placement later.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



10% Discount Coupon Code: ASTR14