How to use StorageGRID

StorageGRID

NetApp
June 13, 2025

Table of Contents

How to use StorageGRID

Explore the Grid Manager

The Grid Manager is the browser-based graphical interface that allows you to configure, manage, and monitor your StorageGRID system.

When you sign in to the Grid Manager, you are connecting to an Admin Node. Each StorageGRID system includes one primary Admin Node and any number of non-primary Admin Nodes. You can connect to any Admin Node, and each Admin Node displays a similar view of the StorageGRID system.

You can access the Grid Manager using a supported web browser.

Grid Manager Dashboard

When you first sign in to the Grid Manager, you can use the Dashboard to monitor system activities at a glance.

The Dashboard includes summary information about system health, storage use, ILM processes, and S3 and Swift operations.
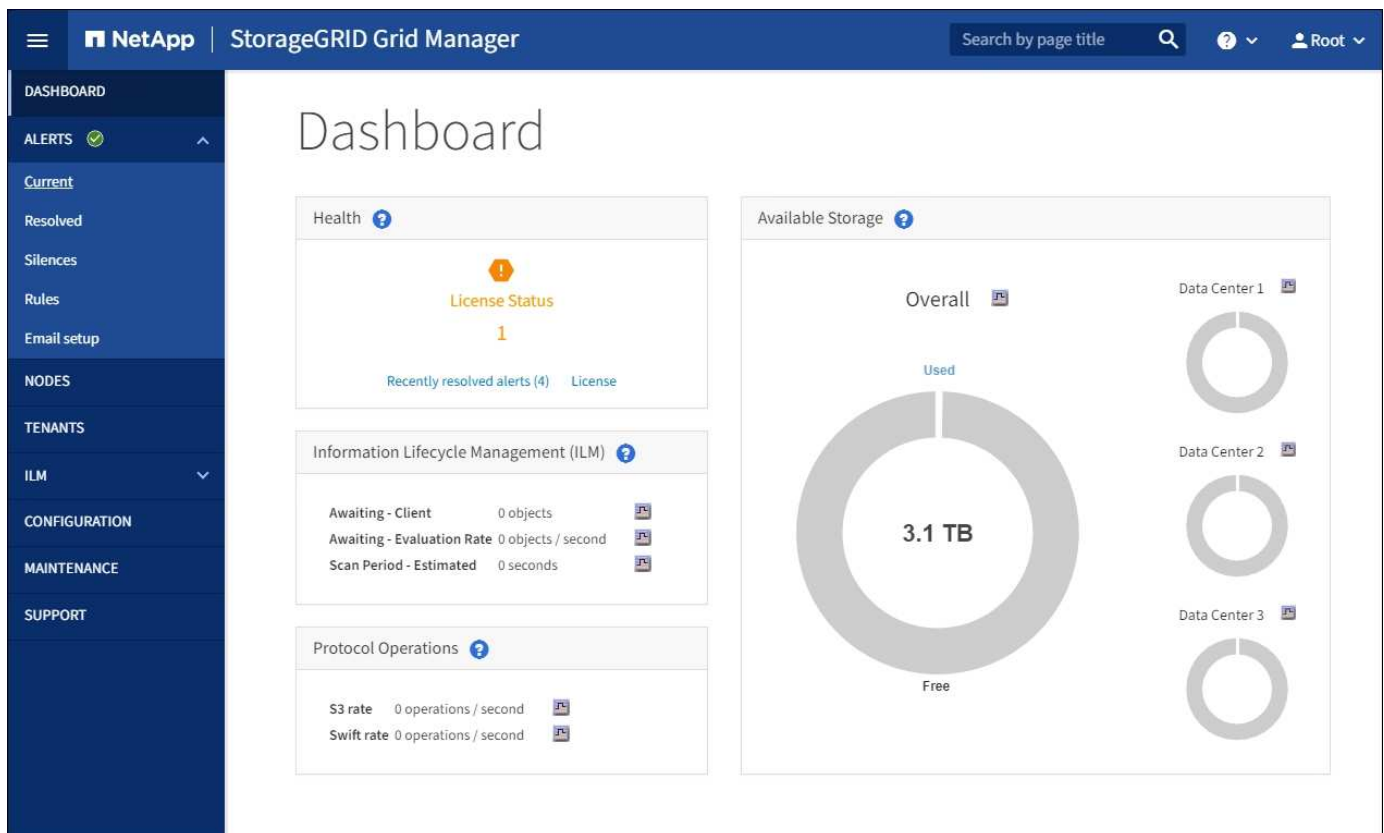
For an explanation of the information on each panel, click the help icon 🅿 for that panel.

Learn more

- Monitor and troubleshoot

Search field

The **Search** field in the header bar allows you to quickly navigate to a specific page within Grid Manager. For example, you can enter **km** to access the Key Management Server (KMS) page. You can use **Search** to find entries in the sidebar of the Grid Manager and on the Configuration, Maintenance, and Support menus.

Alerts menu

The Alerts menu provides an easy-to-use interface for detecting, evaluating, and resolving issues that might occur during StorageGRID operation.
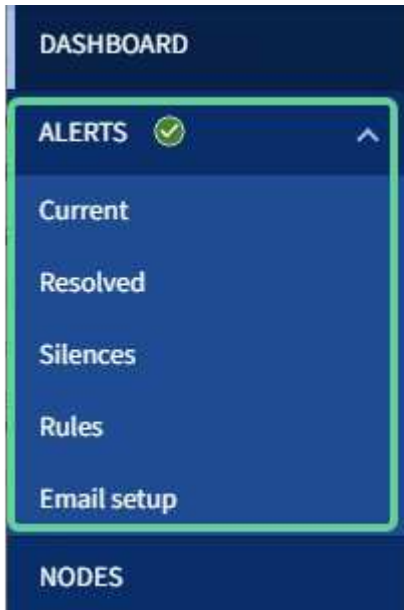
From the Alerts menu, you can do the following:

- Review current alerts
- Review resolved alerts
- Configure silences to suppress alert notifications
- Define alert rules for conditions that trigger alerts
- Configure the email server for alert notifications

Learn more
- Monitor and manage alerts
- Monitor and troubleshoot

Nodes page

The Nodes page displays information about the entire grid, each site in the grid, and each node at a site.

The Nodes home page displays combined metrics for the entire grid. To view information for a particular site or node, select the site or node.

Learn more

- View the Nodes page

- Monitor and troubleshoot

Tenants page

The Tenants page allows you to create and monitor the storage tenant accounts for your StorageGRID system. You must create at least one tenant account to specify who can store and retrieve objects and which functionality is available to them.

The Tenants page also provides usage details for each tenant, including the amount of storage used and the number of objects. If you set a quota when you created the tenant, you can see how much of that quota has been used.

Learn more

- Manage tenants and client connections
- Administer StorageGRID
- Use a tenant account

ILM menu

The ILM menu allows you to configure the information lifecycle management (ILM) rules and policies that govern data durability and availability. You can also enter an object identifier to view the metadata for that object.

Learn more

- Use information lifecycle management
- Manage objects with ILM

Configuration menu

The Configuration menu allows you to specify network settings, security settings, system settings, monitoring options, and access control options.

Configuration

Configure your StorageGRID system.

Network	Security	System	Monitoring	Access control
Domain names	Certificates	Display options	Audit and syslog server	Admin groups
High availability groups	Key management server	Grid options	SNMP agent	Admin users
Link cost	Proxy settings	S3 Object Lock		Grid passwords
Load balancer endpoints	Untrusted Client Networks	Storage options		Identity federation
Traffic classification				Single sign-on
VLAN interfaces				

Learn more

- Configure network settings
- Manage tenants and client connections
- Review audit messages
- Control StorageGRID access
- Administer StorageGRID
- Monitor and troubleshoot
- Review audit logs

Maintenance menu

The Maintenance menu allows you to perform maintenance tasks, system maintenance, and network maintenance.

Tasks

Maintenance tasks include:

- Decommission operations to remove unused grid nodes and sites.
- Expansion operations to add new grid nodes and sites.
- Recovery operations to replace a failed node and restore data.
- Object existence check to verify the existence (although not the correctness) of object data.

System

System maintenance tasks you can perform include:

- Reviewing details for the current StorageGRID license or uploading a new license.
- Generating a Recovery Package.
- Performing StorageGRID software updates, including software upgrades, hotfixes, and updates to the SANtricity OS software on selected appliances.

Network

Network maintenance tasks you can perform include:

- Editing information about DNS servers.
- Configuring the subnets that are used on the Grid Network.
- Editing information about NTP servers.

Learn more
- Perform maintenance
- Download the Recovery Package
- Expand your grid

- [Upgrade software](#)

- [Recover and maintain](#)

- [SG6000 storage appliances](#)

- [SG5700 storage appliances](#)

- [SG5600 storage appliances](#)

Support menu

The Support menu provides options that help technical support analyze and troubleshoot your system. There are two parts to the Support menu: Tools and Alarms (legacy).

Tools

From the Tools section of the Support menu, you can:

- Enable AutoSupport.

- Perform a set of diagnostic checks on the current state of the grid.

- Access the grid topology tree to view detailed information about grid nodes, services, and attributes.

- Retrieve log files and system data.

- Review detailed metrics and charts.

> (i) The tools available from the **Metrics** option are intended for use by technical support. Some features and menu items within these tools are intentionally non-functional.

Alarms (legacy)

From the Alarms (legacy) section of the Support menu, you can review current, historical, and global alarms, set up custom events, and set up email notifications for legacy alarms and AutoSupport.

> ⓘ While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

Learn more

- StorageGRID architecture and network topology
- StorageGRID attributes
- Use StorageGRID support options
- Administer StorageGRID
- Monitor and troubleshoot

Help menu

The Help option provides access to the StorageGRID Documentation Center for the current release and to the API documentation. You can also determine which version of StorageGRID is currently installed.
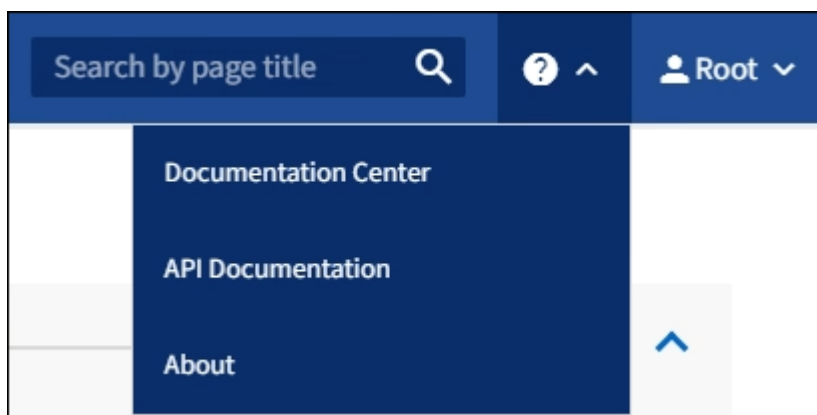
Learn more

- Administer StorageGRID

Explore the Tenant Manager

The Tenant Manager is the browser-based graphical interface that tenant users access to configure, manage, and monitor their storage accounts.

When tenant users sign in to the Tenant Manager, they are connecting to an Admin Node.

Tenant Manager Dashboard

After a grid administrator creates a tenant account using the Grid Manager or the Grid Management API, tenant users can sign in to the Tenant Manager.

The Tenant Manager Dashboard allows tenant users to monitor storage usage at a glance. The Storage usage panel contains a list of the largest buckets (S3) or containers (Swift) for the tenant. The Space used value is the total amount of object data in the bucket or container. The bar chart represents the relative sizes of these buckets or containers.

The value shown above the bar chart is a sum of the space used for all of the tenant's buckets or containers. If the maximum number of gigabytes, terabytes, or petabytes available for the tenant was specified when the

account was created, the amount of quota used and remaining are also shown.

Dashboard

16	**Buckets**	2	**Platform services**	0	**Groups**	1	**User**
	View buckets		endpoints		View groups		View users
			View endpoints				

Storage usage ❓

6.5 TB of 7.2 TB used   0.7 TB (10.1%) remaining

Total objects

8,418,886
objects

Bucket name	Space used	Number of objects
● Bucket-15	969.2 GB	913,425
● Bucket-04	937.2 GB	576,806
● Bucket-13	815.2 GB	957,389
● Bucket-06	812.5 GB	193,843
● Bucket-10	473.9 GB	583,245
● Bucket-03	403.2 GB	981,226
● Bucket-07	362.5 GB	420,726
● Bucket-05	294.4 GB	785,190
● 8 other buckets	1.4 TB	3,007,036

Tenant details ❓

Name: Tenant02
ID: 3341 1240 0546 8283 2208
✓ Platform services enabled
✓ Can use own identity source
✓ S3 Select enabled

Storage menu (S3 tenants only)

The Storage menu is provided for S3 tenant accounts only. This menu allows S3 users to manage access keys, create and delete buckets, and manage platform service endpoints.

DASHBOARD

STORAGE (S3) ︿

My access keys

Buckets

Platform services endpoints

My access keys

S3 tenant users can manage access keys as follows:

- Users who have the Manage Your Own S3 Credentials permission can create or remove their own S3 access keys.

- Users who have the Root Access permission can manage the access keys for the S3 root account, their own account, and all other users. Root access keys also provide full access to the tenant's buckets and objects unless explicitly disabled by a bucket policy.

> (i) Managing the access keys for other users takes place from the Access Management menu.

Buckets

S3 tenant users with the appropriate permissions can perform the following tasks related to buckets:

- Create buckets
- Enable S3 Object Lock for a new bucket (assumes that S3 Object Lock is enabled for the StorageGRID system)
- Update consistency level settings
- Apply a default retention setting
- Configure cross-origin resource sharing (CORS)
- Enable and disable last access time update settings for the buckets belonging to the tenant
- Delete empty buckets
- Manage the objects in a bucket using the experimental S3 Console

If a grid administrator has enabled the use of platform services for the tenant account, an S3 tenant user with the appropriate permissions can also perform these tasks:

- Configure S3 event notifications, which can be sent to a destination service that supports the AWS Simple Notification Service™ (SNS).
- Configure CloudMirror replication, which enables the tenant to automatically replicate objects to an external S3 bucket.
- Configure search integration, which sends object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

Platform services endpoints

If a grid administrator has enabled the use of platform services for the tenant account, an S3 tenant user with the Manage Endpoints permission can configure a destination endpoint for each platform service.

Access Management menu

The Access Management menu allows StorageGRID tenants to import user groups from a federated identity source and assign management permissions. Tenants can also manage local tenant groups and users, unless single sign-on (SSO) is in effect for the entire StorageGRID system.
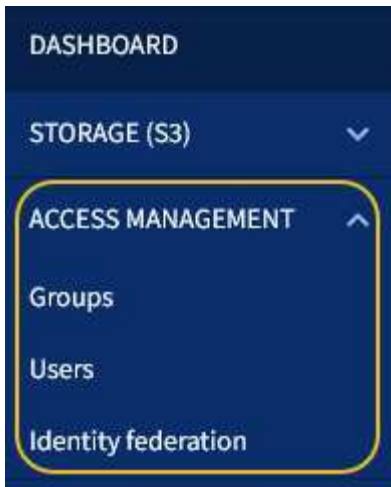
Related information

- Explore the Grid Manager
- Use a tenant account

Control StorageGRID access

You control who can access StorageGRID and which tasks users can perform by creating or importing groups and users and assigning permissions to each group. Optionally, you can enable single sign-on (SSO), create client certificates, and change grid passwords.

Control access to the Grid Manager

You determine who can access the Grid Manager and the Grid Management API by importing groups and users from an identity federation service or by setting up local groups and local users.

Using identity federation makes setting up groups and users faster, and it allows users to sign in to StorageGRID using familiar credentials. You can configure identity federation if you use Active Directory, OpenLDAP, or Oracle Directory Server.

> ℹ️ Contact technical support if you want to use another LDAP v3 service.

You determine which tasks each user can perform by assigning different permissions to each group. For example, you might want users in one group to be able to manage ILM rules and users in another group to perform maintenance tasks. A user must belong to at least one group to access the system.

Optionally, you can configure a group to be read-only. Users in a read-only group can only view settings and features. They cannot make any changes or perform any operations in the Grid Manager or Grid Management API.

Enable single sign-on

The StorageGRID system supports single sign-on (SSO) using the Security Assertion Markup Language 2.0 (SAML 2.0) standard. When SSO is enabled, all users must be authenticated by an external identity provider before they can access the Grid Manager, the Tenant Manager, the Grid Management API, or the Tenant Management API. Local users cannot sign in to StorageGRID.

When SSO is enabled and users sign in to StorageGRID, they are redirected to your organization's SSO page

to validate their credentials. When users sign out of one Admin Node, they are automatically signed out of all Admin Nodes.

Change grid passwords

The provisioning passphrase is required for many installation and maintenance procedures, and for downloading the StorageGRID Recovery Package. The passphrase is also required to download backups of the grid topology information and encryption keys for the StorageGRID system. You can change this passphrase as required.

Related information

- Administer StorageGRID
- Use a tenant account

Manage tenants and client connections

As a grid administrator, you create and manage the tenant accounts that S3 and Swift clients use to store and retrieve objects, and manage the configuration options that control how clients connect to your StorageGRID system.

Tenant accounts

A tenant account allows you to specify who can use your StorageGRID system to store and retrieve objects, and which functionality is available to them. Tenant accounts allow client applications that support the S3 REST API or the Swift REST API to store and retrieve objects on StorageGRID. Each tenant account uses either the S3 client protocol or the Swift client protocol.

You must create at least one tenant account for each client protocol that will be used to store objects on your StorageGRID system. Optionally, you can create additional tenant accounts if you want to segregate the objects stored on your system by different entities. Each tenant account has its own federated or local groups and users, and its own buckets (containers for Swift) and objects.

You can use the Grid Manager or the Grid Management API to create tenant accounts. When creating a tenant account, you specify the following information:

- Display name for the tenant (the tenant's account ID is assigned automatically and cannot be changed).
- Whether the tenant account will use the S3 or Swift.
- For S3 tenant accounts: Whether the tenant account is allowed to use platform services. If the use of platform services is allowed, the grid must be configured to support their use.
- Optionally, a storage quota for the tenant account—the maximum number of gigabytes, terabytes, or petabytes available for the tenant's objects. A tenant's storage quota represents a logical amount (object size), not a physical amount (size on disk).
- If identity federation is enabled for the StorageGRID system, which federated group has Root Access permission to configure the tenant account.
- If single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid's identity source, and the initial password for the tenant's local root user.

If S3 tenant accounts need to comply with regulatory requirements, grid administrators can enable the global S3 Object Lock setting for the StorageGRID system. When S3 Object Lock is enabled for the system, all S3

tenant accounts can create buckets with S3 Object Lock enabled and then specify retention and legal hold settings for the object versions in that bucket.

After a tenant account is created, tenant users can sign in to the Tenant Manager.

Client connections to StorageGRID nodes

Before tenant users can use S3 or Swift clients to store and retrieve data in StorageGRID, you must decide how these clients will connect to StorageGRID nodes.

Client applications can store or retrieve objects by connecting to any of the following:

- The Load Balancer service on Admin Nodes or Gateway Nodes. This is the recommended connection.
- The CLB service on Gateway Nodes.

  > ⓘ The CLB service is deprecated.

- Storage Nodes, with or without an external load balancer.

When configuring StorageGRID so that clients can use the Load Balancer service, you perform the following steps:

1. Optionally configure high availability (HA) groups. If you create an HA group, the interfaces of multiple Admin Nodes and Gateway Nodes are placed into an active-backup configuration. Client connections are made using the virtual IP address of the HA group.

2. Configure endpoints for the Load Balancer service. The Load Balancer service on Admin Nodes or Gateway Nodes distributes incoming network connections from client applications to Storage Nodes. When creating a load balancer endpoint, you specify a port number, whether the endpoint accepts HTTP or HTTPS connections, the type of client (S3 or Swift) that will use the endpoint, and the certificate to be used for HTTPS connections (if applicable).

3. Optionally specify that a node's Client Network is untrusted to ensure that all connections to the node's Client Network occur on the load balancer endpoints.

Related information

- Administer StorageGRID
- Use a tenant account
- Use S3
- Use Swift
- Explore the Tenant Manager
- Configure network settings

Configure network settings

You can configure various network settings from the Grid Manager to fine tune the operation of your StorageGRID system.

Domain names

If you plan to support S3 virtual hosted-style requests, you must configure the list of endpoint domain names that S3 clients connect to. Examples include `s3.example.com`, `s3.example.co.uk`, and `s3-east.example.com`.

The configured server certificates must match the endpoint domain names.

High availability groups

You can use high availability (HA) groups to provide highly available data connections for S3 and Swift clients or to provide highly available connections to the Grid Manager and the Tenant Manager.

When you create an HA group, you select a network interface for one or more nodes. Each HA group provides access to the shared services on the selected nodes.

- HA groups that include interfaces on Gateway Nodes, Admin Nodes, or both provide highly available data connections for S3 and Swift clients.
- HA groups that include interfaces on Admin Nodes only provide highly available connections to the Grid Manager and the Tenant Manager.

The interfaces can belong to the Grid Network (eth0), the Client Network (eth2), or a VLAN network.

You can assign up to 10 virtual IP (VIP) addresses to each HA group. You specify one interface to be the Primary interface and rank any other interfaces in priority order. The Primary interface is the active interface unless a failure occurs. If the active interface fails, the VIP addresses move to the first backup interface in the priority order. If that interface fails, the VIP addresses move to the next backup interface, and so on.

Link costs

You can adjust link costs to reflect the latency between sites. When two or more data center sites exist, link costs prioritize which data center site should provide a requested service.

Load balancer endpoints

You can use a load balancer to handle ingest and retrieval workloads from S3 and Swift clients. Load balancing maximizes speed and connection capacity by distributing the workloads and connections across multiple Storage Nodes.

If you want to use the StorageGRID load balancer service, which is included on Admin Nodes and Gateway Nodes, you must configure one or more load balancer endpoints. Each endpoint defines a Gateway Node or Admin Node port for S3 and Swift requests to Storage Nodes.

Traffic classification

You can create traffic classification policies to identify and handle different types of network traffic, including traffic related to specific buckets, tenants, client subnets, or load balancer endpoints. These policies can assist with traffic limiting and monitoring.

VLAN interfaces

You can create virtual LAN (VLAN) interfaces to isolate and partition traffic for security, flexibility, and performance. Each VLAN interface is associated with one or more parent interfaces on Admin Nodes and

Gateway Nodes. You can use VLAN interfaces in HA groups and in load balancer endpoints to segregate client or admin traffic by application or tenant.

For example, your network might use VLAN 100 for FabricPool traffic and VLAN 200 for an archive application.

Related information

- Administer StorageGRID

- Manage tenants and client connections

Configure security settings

You can configure various security settings from the Grid Manager to help secure your StorageGRID system.

Certificates

StorageGRID uses two types of security certificates:

- Server certificates are required when you use HTTPS connections. Server certificates are used to establish secure connections between clients and servers, authenticating the identity of a server to its clients and providing a secure communication path for data. The server and the client each have a copy of the certificate.

- Client certificates authenticate a client or user identity to the server, providing more secure authentication than passwords alone. Client certificates do not encrypt data.

You can view all StorageGRID certificates on the **CONFIGURATION** > **Security** > **Certificates** page.

Key management servers

You can configure one or more external key management servers (KMS) to provide encryption keys to StorageGRID services and storage appliances. Each KMS or KMS cluster uses the Key Management Interoperability Protocol (KMIP) to provide an encryption key to the appliance nodes at the associated StorageGRID site. Using key management servers lets you protect StorageGRID data even if an appliance is removed from the data center. After the appliance volumes are encrypted, you cannot access any data on the appliance unless the node can communicate with the KMS.

> ⓘ To use encryption key management, you must enable the **Node Encryption** setting for each appliance during installation, before the appliance is added to the grid.

Proxy settings

If you are using S3 platform services or Cloud Storage Pools, you can configure a non-transparent proxy server between Storage Nodes and the external S3 endpoints. If you send AutoSupport messages using HTTPS or HTTP, you can configure a non-transparent proxy server between Admin Nodes and technical support.

Proxy Settings
Storage
Admin

Untrusted Client Networks

If you are using a Client Network, you can help secure StorageGRID from hostile attacks by specifying that the Client Network on each node be untrusted. If a node's Client Network is untrusted, the node only accepts inbound connections on ports explicitly configured as load balancer endpoints.

For example, you might want a Gateway Node to refuse all inbound traffic on the Client Network except for HTTPS S3 requests. Or, you might want to enable outbound S3 platform service traffic from a Storage Node, while preventing any inbound connections to that Storage Node on the Client Network.

Related information

- Administer StorageGRID
- Manage tenants and client connections

Configure system settings

You can configure various system settings from the Grid Manager to fine tune the operation of your StorageGRID system.

Display options

Display options allow you to specify the timeout period for user sessions and to suppress email notifications for legacy alarms and event-triggered AutoSupport messages.

Grid options

You can use Grid Options to configure the settings for all of the objects stored in your StorageGRID system, including stored object compression, stored object encryption. and stored object hashing.

You can also use these options to specify global settings for S3 and Swift client operations.

S3 Object Lock

The StorageGRID S3 Object Lock feature is an object-protection solution that is equivalent to S3 Object Lock in Amazon Simple Storage Service (Amazon S3). You can enable the global S3 Object Lock setting for a StorageGRID system to allow S3 tenant accounts to create buckets with S3 Object Lock enabled. The tenant can then use an S3 client application to optionally specify retention settings (retain until date, legal hold, or both) for the objects in those buckets. In addition, each bucket that has S3 Object Lock enabled can optionally have a default retention mode and retention period, which apply if objects are added to the bucket without their own retention settings.
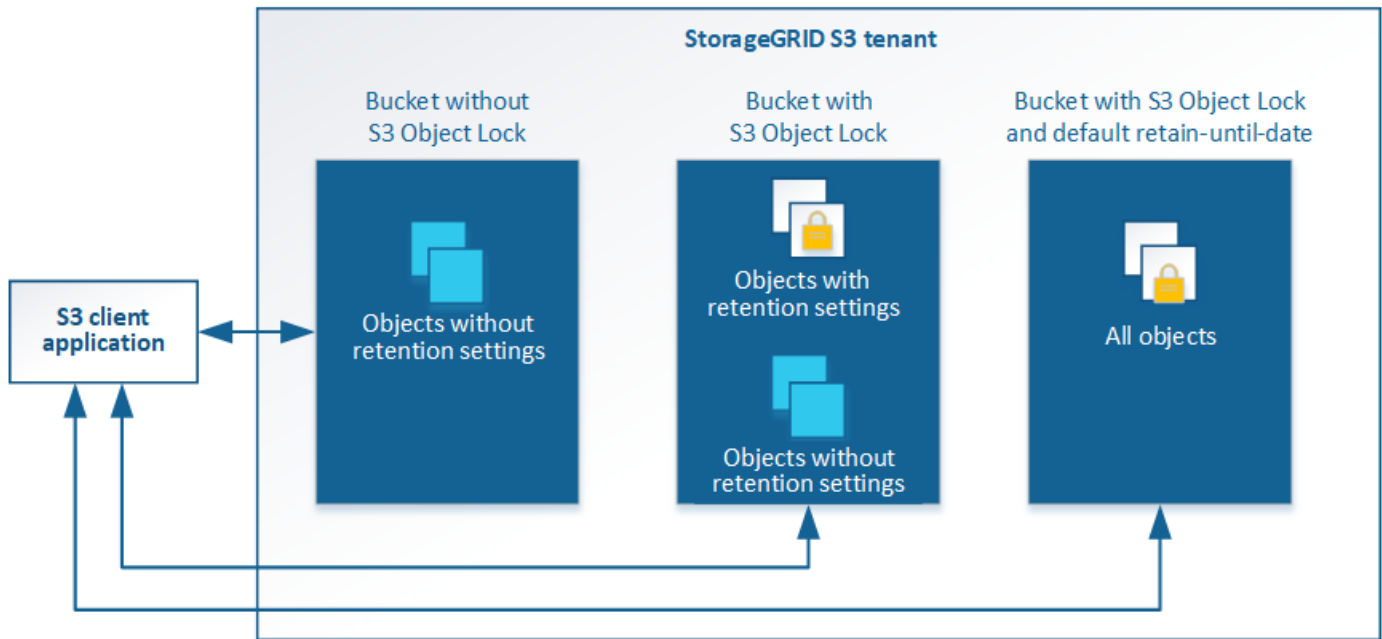
Storage options

Storage options allow you to control object segmentation and to override storage volume watermark settings to manage a Storage Node's usable storage space.

Use information lifecycle management

You use information lifecycle management (ILM) to control the placement, duration, and data protection for all objects in your StorageGRID system. ILM rules determine how StorageGRID stores objects over time. You configure one or more ILM rules and then add them to an ILM policy.

ILM rules define:

- Which objects should be stored. A rule can apply to all objects, or you can specify filters to identify which objects a rule applies to. For example, a rule can apply only to objects associated with certain tenant accounts, specific S3 buckets or Swift containers, or specific metadata values.
- The storage type and location. Objects can be stored on Storage Nodes, in Cloud Storage Pools, or on Archive Nodes.
- The type of object copies made. Copies can be replicated or erasure coded.
- For replicated copies, the number of copies made.
- For erasure coded copies, the erasure-coding scheme used.
- The changes over time to an object's storage location and type of copies.
- How object data is protected as objects are ingested into the grid (synchronous placement or dual commit).

Note that object metadata is not managed by ILM rules. Instead, object metadata is stored in a Cassandra database in what is known as a metadata store. Three copies of object metadata are automatically maintained at each site to protect the data from loss. The copies are evenly distributed across all Storage Nodes.

Example ILM rule

This example ILM rule applies to the objects belonging to Tenant A. It makes two replicated copies of those objects and stores each copy at a different site. The two copies are retained "forever," which means that StorageGRID will not automatically delete them. Instead, StorageGRID will retain these objects until they are deleted by a client delete request or by the expiration of a bucket lifecycle.

This rule uses the Balanced option for ingest behavior: the two-site placement instruction is applied as soon as Tenant A saves an object to StorageGRID, unless it is not possible to immediately make both required copies. For example, if Site 2 is unreachable when Tenant A saves an object, StorageGRID will make two interim copies on Storage Nodes at Site 1. As soon as Site 2 becomes available, StorageGRID will make the required copy at that site.

How an ILM policy evaluates objects

The active ILM policy for your StorageGRID system controls the placement, duration, and data protection of all objects.

When clients save objects to StorageGRID, the objects are evaluated against the ordered set of ILM rules in the active policy, as follows:

1. If the filters for the first rule in the policy match an object, the object is ingested according to that rule's ingest behavior and stored according to that rule's placement instructions.

2. If the filters for the first rule do not match the object, the object is evaluated against each subsequent rule in the policy until a match is made.

3. If no rules match an object, the ingest behavior and placement instructions for the default rule in the policy are applied. The default rule is the last rule in a policy and cannot use any filters. It must apply to all tenants, all buckets, and all object versions.

Example ILM policy

This example ILM policy uses three ILM rules.

In this example, Rule 1 matches all objects belonging to Tenant A. These objects are stored as three replicated copies at three sites. Objects belonging to other tenants are not matched by Rule 1, so they are evaluated against Rule 2.

Rule 2 matches all objects from other tenants but only if they are greater than 1 MB. These larger objects are stored using 6+3 erasure coding at three sites. Rule 2 does not match objects 1 MB or smaller, so these objects are evaluated against Rule 3.

Rule 3 is the last and default rule in the policy, and it does not use filters. Rule 3 makes two replicated copies of all objects not matched by Rule 1 or Rule 2 (objects not belonging to Tenant A that are 1 MB or smaller).
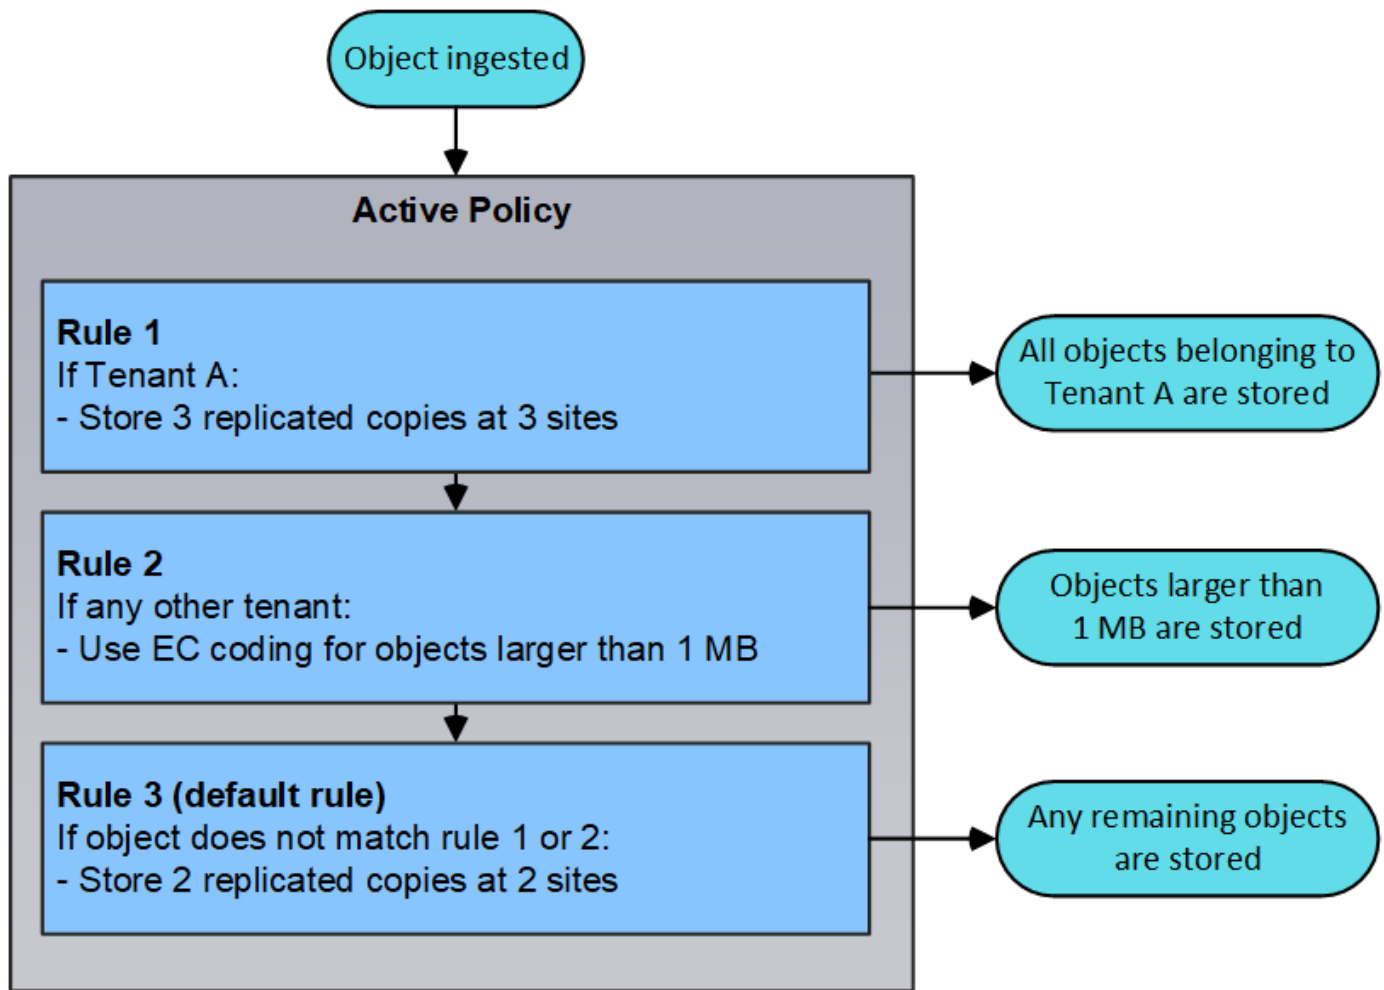
Related information

- Manage objects with ILM

Monitor operations

View the Nodes page

When you need more detailed information about your StorageGRID system than the Dashboard provides, you can use the Nodes page to view metrics for the entire grid, each site in the grid, and each node at a site.

The Nodes table lists all the sites and nodes in your StorageGRID system. Summary information is displayed for each node. If a node has an active alert, an icon appears next to the node name. If the node is connected and has no active alerts, no icon is shown.

Connection state icons

- **Not connected - Unknown** : The node is not connected to the grid for an unknown reason. For example, the network connection between nodes has been lost or the power is down. The **Unable to communicate with node** alert might also be triggered. Other alerts might be active as well. This situation requires immediate attention.

  > A node might appear as Unknown during managed shutdown operations. You can ignore the Unknown state in these cases.

- **Not connected - Administratively down** : The node is not connected to the grid for an expected reason. For example, the node, or services on the node, has been gracefully shut down, the node is rebooting, or the software is being upgraded. One or more alerts might also be active.

If a node is disconnected from the grid, it might have an underlying alert, but only the "Not connected" icon appears. To see the active alerts for a node, select the node.

Alert icons

If there is an active alert for a node, one of the following icons appears next to the node name:

- **Critical** : An abnormal condition exists that has stopped the normal operations of a StorageGRID node

or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved.

- **Major** ⚠️ : An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service.

- **Minor** ⚠️ : The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that do not clear on their own to ensure they do not result in a more serious problem.

Details for a system, site, or node

To view the available information, select the name of the grid, site, or node as follows:

- Select the grid name to see an aggregate summary of the statistics for your entire StorageGRID system. (The screenshot shows a system named StorageGRID Deployment.)
- Select a specific data center site to see an aggregate summary of the statistics for all nodes at that site.
- Select a specific node to view detailed information for that node.

Tabs for the Nodes page

The tabs at the top of the Nodes page are based on what you select from the tree at the left.

Tab name	Description	Included for
Overview	• Provides basic information about each node. • Shows any active alerts affecting the node.	All nodes
Hardware	• Displays CPU utilization and memory usage for each node • For appliance nodes, provides additional hardware information.	All nodes
Network	Displays a graph showing the network traffic received and sent across the network interfaces. The view for a single node shows additional information for the node.	All nodes, each site, and the entire grid
Storage	• Provides details for the disk devices and volumes on each node. • For Storage Nodes, each site, and the entire grid, includes graphs showing object data storage and metadata storage used over time.	All nodes, each site, and the entire grid
Objects	• Provides information about S3 and Swift ingest and retrieve rates. • For Storage Nodes, provides object counts and information about metadata store queries and background verification.	Storage Nodes, each site, and the entire grid

Tab name	Description	Included for
ILM	Provides information about Information Lifecycle Management (ILM) operations. • For Storage Nodes, provides details about ILM evaluation and background verification for erasure coded objects. • For each site and the entire grid, shows a graph of the ILM queue over time. • For the entire grid, provides the estimated time to complete a full ILM scan of all objects.	Storage Nodes, each site, and the entire grid
Load balancer	Includes performance and diagnostic graphs related to the Load Balancer service. • For each site, provides an aggregate summary of the statistics for all nodes at that site. • For the entire grid, provides an aggregate summary of the statistics for all sites.	Admin Nodes and Gateway Nodes, each site, and the entire grid
Platform services	Provides information about any S3 platform service operations at a site.	Each site
SANtricity System Manager	Provides access to SANtricity System Manager. From SANtricity System Manager, you can review hardware diagnostic and environmental information for the storage controller, as well as issues related to the drives.	Storage appliance nodes **Note**: The SANtricity System Manager tab will not appear if the controller firmware on the storage appliance is earlier than 8.70 (11.70).

Prometheus metrics

The Prometheus service on Admin Nodes collects time series metrics from the services on all nodes.

The metrics collected by Prometheus are used in a number of places in the Grid Manager:

• **Nodes page**: The graphs and charts on the tabs available from the Nodes page use the Grafana visualization tool to display the time-series metrics collected by Prometheus. Grafana displays time-series data in graph and chart formats, while Prometheus serves as the backend data source.

- **Alerts**: Alerts are triggered at specific severity levels when alert rule conditions that use Prometheus metrics evaluate as true.

- **Grid Management API**: You can use Prometheus metrics in custom alert rules or with external automation tools to monitor your StorageGRID system. A complete list of Prometheus metrics is available from the Grid Management API. (From the top of the Grid Manager, select the help icon and select **API Documentation** > **metrics**.) While more than a thousand metrics are available, only a relatively small number are required to monitor the most critical StorageGRID operations.

> ℹ️ Metrics that include *private* in their names are intended for internal use only and are subject to change between StorageGRID releases without notice.

- The **SUPPORT** > **Tools** > **Diagnostics** page and the **SUPPORT** > **Tools** > **Metrics** page: These pages, which are primarily intended for use by technical support, provide a number of tools and charts that use the values of Prometheus metrics.

> ℹ️ Some features and menu items within the Metrics page are intentionally non-functional and are subject to change.

StorageGRID attributes

Attributes report values and statuses for many of the functions of the StorageGRID system. Attribute values are available for each grid node, each site, and the entire grid.

StorageGRID attributes are used in a number of places in the Grid Manager:

- **Nodes page**: Many of the values shown on the Nodes page are StorageGRID attributes. (Prometheus metrics are also shown on the Nodes pages.)

- **Alarms**: When attributes reach defined threshold values, StorageGRID alarms (legacy system) are triggered at specific severity levels.

- **Grid Topology tree**: Attribute values are shown in the Grid Topology tree (**SUPPORT** > **Tools** > **Grid topology**).

- **Events**: System events occur when certain attributes record an error or fault condition for a node, including errors such as network errors.

Attribute values

Attributes are reported on a best-effort basis and are approximately correct. Attribute updates can be lost under some circumstances, such as the crash of a service or the failure and rebuild of a grid node.

In addition, propagation delays might slow the reporting of attributes. Updated values for most attributes are sent to the StorageGRID system at fixed intervals. It can take several minutes before an update is visible in the system, and two attributes that change more or less simultaneously can be reported at slightly different times.

Related information

- Monitor and troubleshoot
- Monitor and manage alerts
- Use StorageGRID support options

Monitor and manage alerts

The alert system provides an easy-to-use interface for detecting, evaluating, and resolving the issues that can occur during StorageGRID operation.

The alert system is designed to be your primary tool for monitoring any issues that might occur in your StorageGRID system.

- The alert system focuses on actionable problems in the system. Alerts are triggered for events that require your immediate attention, not for events that can safely be ignored.

- The Current Alerts and Resolved Alerts pages provide a user friendly interface for viewing current and historical problems. You can sort the listing by individual alerts and alert groups. For example, you might want to sort all alerts by node/site to see which alerts are affecting a specific node. Or, you might want to sort the alerts in a group by time triggered to find the most recent instance of a specific alert.

- Multiple alerts of the same type are grouped into one email to reduce the number of notifications. In addition, multiple alerts of the same type are shown as a group on the Current Alerts and Resolved Alerts pages. You can expand and collapse alert groups to show or hide the individual alerts. For example, if several nodes are reporting the **Unable to communicate with node** alert, only one email is sent and the alert is shown as a group on the Current Alerts page.

- Alerts use intuitive names and descriptions to help you understand more quickly what the problem is. Alert notifications include details about the node and site affected, the alert severity, the time when the alert rule was triggered, and the current value of metrics related to the alert.

- Alert email notifications and the alert listings on the Current Alerts and Resolved Alerts pages provide recommended actions for resolving an alert. These recommended actions often include direct links to StorageGRID documentation to make it easier to find and access more detailed troubleshooting procedures.

Low installed node memory

The amount of installed memory on a node is low.

Recommended actions

Increase the amount of RAM available to the virtual machine or Linux host. Check the threshold value for the major alert to determine the default minimum requirement for a StorageGRID node.

See the instructions for your platform:

- VMware installation
- Red Hat Enterprise Linux or CentOS installation
- Ubuntu or Debian installation

Time triggered

2019-07-15 17:07:41 MDT *(2019-07-15 23:07:41 UTC)*

Status
Active (silence this alert ⬈)

Site / Node
Data Center 2 / DC2-S1-99-56

Severity
❌ Critical

Total RAM size
8.38 GB

Condition
View conditions | Edit rule ⬈

Close

ⓘ The legacy alarm system is deprecated. The user interface and APIs for the legacy alarm system will be removed in a future release. The alert system offers significant benefits and is easier to use.

Manage alerts

All StorageGRID users can view alerts. If you have the Root Access or Manage Alerts permission, you can also manage alerts, as follows:

- If you need to temporarily suppress the notifications for an alert at one or more severity levels, you can easily silence a specific alert rule for a specified duration. You can silence an alert rule for the entire grid, a single site, or a single node.

- You can edit the default alert rules as required. You can disable an alert rule completely, or change its trigger conditions and duration.

- You can create custom alert rules to target the specific conditions that are relevant to your situation and to provide your own recommended actions. To define the conditions for a custom alert, you create expressions using the Prometheus metrics available from the Metrics section of the Grid Management API.

  For example, this expression causes an alert to be triggered if the amount of installed RAM for a node is less than 24,000,000,000 bytes (24 GB).

```
node_memory_MemTotal < 24000000000
```

Related information

Use SNMP monitoring

If you want to monitor StorageGRID using the Simple Network Management Protocol (SNMP), you can use the Grid Manager to configure the SNMP agent.

Each StorageGRID node runs an SNMP agent, or daemon, that provides a management information base (MIB). The StorageGRID MIB contains table and notification definitions for alerts and alarms. Each StorageGRID node also supports a subset of MIB-II objects.

Initially, SNMP is disabled on all nodes. When you configure the SNMP agent, all StorageGRID nodes receive the same configuration.

The StorageGRID SNMP agent supports all three versions of the SNMP protocol. The agent provides read-only MIB access for queries, and it can send two types of event-driven notifications to a management system:

- **Traps** are notifications sent by the SNMP agent that do not require acknowledgment by the management system. Traps serve to notify the management system that something has happened within StorageGRID, such as an alert being triggered. Traps are supported in all three versions of SNMP.

- **Informs** are similar to traps, but they require acknowledgment by the management system. If the SNMP agent does not receive an acknowledgment within a certain amount of time, it resends the inform until an acknowledgment is received or the maximum retry value has been reached. Informs are supported in SNMPv2c and SNMPv3.

Trap and inform notifications are sent in the following cases:

- A default or custom alert is triggered at any severity level. To suppress SNMP notifications for an alert, you must configure a silence for the alert. Alert notifications are sent by whichever Admin Node is configured to be the preferred sender.

- Certain alarms (legacy system) are triggered at specified severity levels or higher.

> (i) SNMP notifications are not sent for every alarm or every alarm severity.

Related information

- Monitor and troubleshoot

Review audit messages

Audit messages can help you get a better understanding of the detailed operations of your StorageGRID system. You can use audit logs to troubleshoot issues and to evaluate performance.

During normal system operation, all StorageGRID services generate audit messages, as follows:

- System audit messages are related to the auditing system itself, grid node states, system-wide task activity, and service backup operations.

- Object storage audit messages are related to the storage and management of objects within StorageGRID, including object storage and retrievals, grid-node to grid-node transfers, and verifications.

- Client read and write audit messages are logged when an S3 or Swift client application makes a request to create, modify, or retrieve an object.

- Management audit messages log user requests to the Management API.

Each Admin Node stores audit messages in text files. The audit share contains the active file (audit.log) as well as compressed audit logs from previous days. Additionally, each node in your grid stores a limited amount of audit messages in a local log file (localaudit.log).

For easy access to audit logs, you can configure client access to the audit share for both NFS and CIFS (CIFS is deprecated). You can also access audit log files directly from the command line of the Admin Node.

Optionally, you can send audit information stored on Admin Nodes and local nodes to an external syslog server. Using an external syslog server can make it easier to manage your audit information and reduce network traffic. See Configure audit messages and log destinations for more information.

For details on the audit log file, the format of audit messages, the types of audit messages, and the tools available to analyze audit messages, see the instructions for audit messages. To learn how to configure audit client access, see Configure audit client access.

Related information

- Review audit logs

- Administer StorageGRID

Perform maintenance

You perform various maintenance procedures to keep your StorageGRID system up-to-date and to ensure it is performing efficiently. The Grid Manager provides tools and options to facilitate the process of performing maintenance tasks.

Software updates

You can perform three types of software updates from the Software Update page in the Grid Manager:

- StorageGRID software upgrade

- StorageGRID hotfix

- SANtricity OS upgrade

StorageGRID software upgrades

When a new StorageGRID feature release is available, the Software Upgrade page guides you through the process of uploading the required file and upgrading your StorageGRID system. You must upgrade all grid nodes for all data center sites from the primary Admin Node.

During a StorageGRID software upgrade, client applications can continue to ingest and retrieve object data.

Hotfixes

If issues with the software are detected and resolved between feature releases, you might need to apply a hotfix to your StorageGRID system.

StorageGRID hotfixes contain software changes that are made available outside of a feature or patch release.

The same changes are included in a future release.

The StorageGRID Hotfix page, shown below, allows you to upload a hotfix file.

The hotfix is applied first to the primary Admin Node. Then, you must approve the application of the hotfix to other grid nodes until all nodes in your StorageGRID system are running the same software version. You can customize the approval sequence by selecting to approve individual grid nodes, groups of grid nodes, or all grid nodes.

> ⓘ While all grid nodes are updated with the new hotfix version, the actual changes in a hotfix might only affect specific services on specific types of nodes. For example, a hotfix might only affect the LDR service on Storage Nodes.

SANtricity OS upgrades

You might need to upgrade the SANtricity OS Software on the storage controllers of your storage appliances, if the controllers are not functioning optimally. You can upload the SANtricity OS file to the primary Admin Node in your StorageGRID system and apply the upgrade from the Grid Manager.

The SANtricity page, shown below, allows you to upload the SANtricity OS upgrade file.

SANtricity OS

Use this procedure to upgrade the SANtricity OS software (controller firmware) on the storage controllers in your storage appliances.

1. Download the SANtricity OS version that is compatible with the storage controllers. If you use different appliance models, repeat these steps for each model.
2. Confirm the storage controllers are Nominal (**NODES > *appliance node* > Hardware**) and ready to upgrade.
3. Start the upgrade and approve the nodes you want to upgrade. Nodes are upgraded one at a time.
   During the upgrade, a health check is performed and valid NVSRAM is installed. When the upgrade is complete, the appliance is rebooted. The upgrade can take up to 30 minutes for each appliance.
4. Select **Skip Nodes and Finish** if you only want to apply this upgrade to some nodes or if you want to upgrade some nodes later.

SANtricity OS Upgrade File

SANtricity OS Upgrade File ❔   [Browse]

Passphrase

Provisioning Passphrase ❔   [ ] [Start]

After you upload the file, you can approve the upgrade on individual Storage Nodes or all nodes. The ability to selectively approve nodes makes it easier for you to schedule the upgrade. After you approve a node for upgrade, the system performs a health check and installs the upgrade if it is applicable to the node.

Expansion procedures

You can expand a StorageGRID system by adding storage volumes to Storage Nodes, adding new grid nodes to an existing site, or adding a new data center site. If you have Storage Nodes that use the SG6060 or SG6060X storage appliance, you can add one or two expansion shelves to double or triple the storage capacity of the node.

You can perform expansions without interrupting the operation of your current system. When you add nodes or a site, you first deploy the new nodes and then perform the expansion procedure from the Grid Expansion page.

Grid Expansion

ⓘ A new Recovery Package has been generated as a result of the configuration change. Go to the Recovery Package page to download it.

Expansion Progress

Lists the status of grid configuration tasks required to change the grid topology. These grid configuration tasks are run automatically by the StorageGRID system.

| 1. Installing Grid Nodes | | | | In Progress |

Grid Node Status

Lists the installation and configuration status of each grid node included in the expansion.

Search 🔍

Name	Site	Grid Network IPv4 Address	Progress	Stage
DC2-ADM1-184	Site A	172.17.3.184/21		Waiting for NTP to synchronize
DC2-S1-185	Site A	172.17.3.185/21		Waiting for Dynamic IP Service peers
DC2-S2-186	Site A	172.17.3.186/21		Waiting for NTP to synchronize
DC2-S3-187	Site A	172.17.3.187/21		Waiting for NTP to synchronize
DC2-S4-188	Site A	172.17.3.188/21		Waiting for Dynamic IP Service peers
DC2-ARC1-189	Site A	172.17.3.189/21		Waiting for NTP to synchronize

2. Initial Configuration	Pending
3. Distributing the new grid node's certificates to the StorageGRID system.	Pending
4. Starting services on the new grid nodes	Pending
5. Cleaning up unused Cassandra keys	Pending

Node recovery procedures

Grid nodes can fail if a hardware, virtualization, operating system, or software fault renders the node inoperable or unreliable.

The steps to recover a grid node depend on the platform where the grid node is hosted and on the type of grid node. Each type of grid node has a specific recovery procedure, which you must follow exactly. Generally, you try to preserve data from the failed grid node where possible, repair or replace the failed node, use the Recovery page to configure the replacement node, and restore the node's data.

For example, this flowchart shows the recovery procedure if an Admin Node has failed.
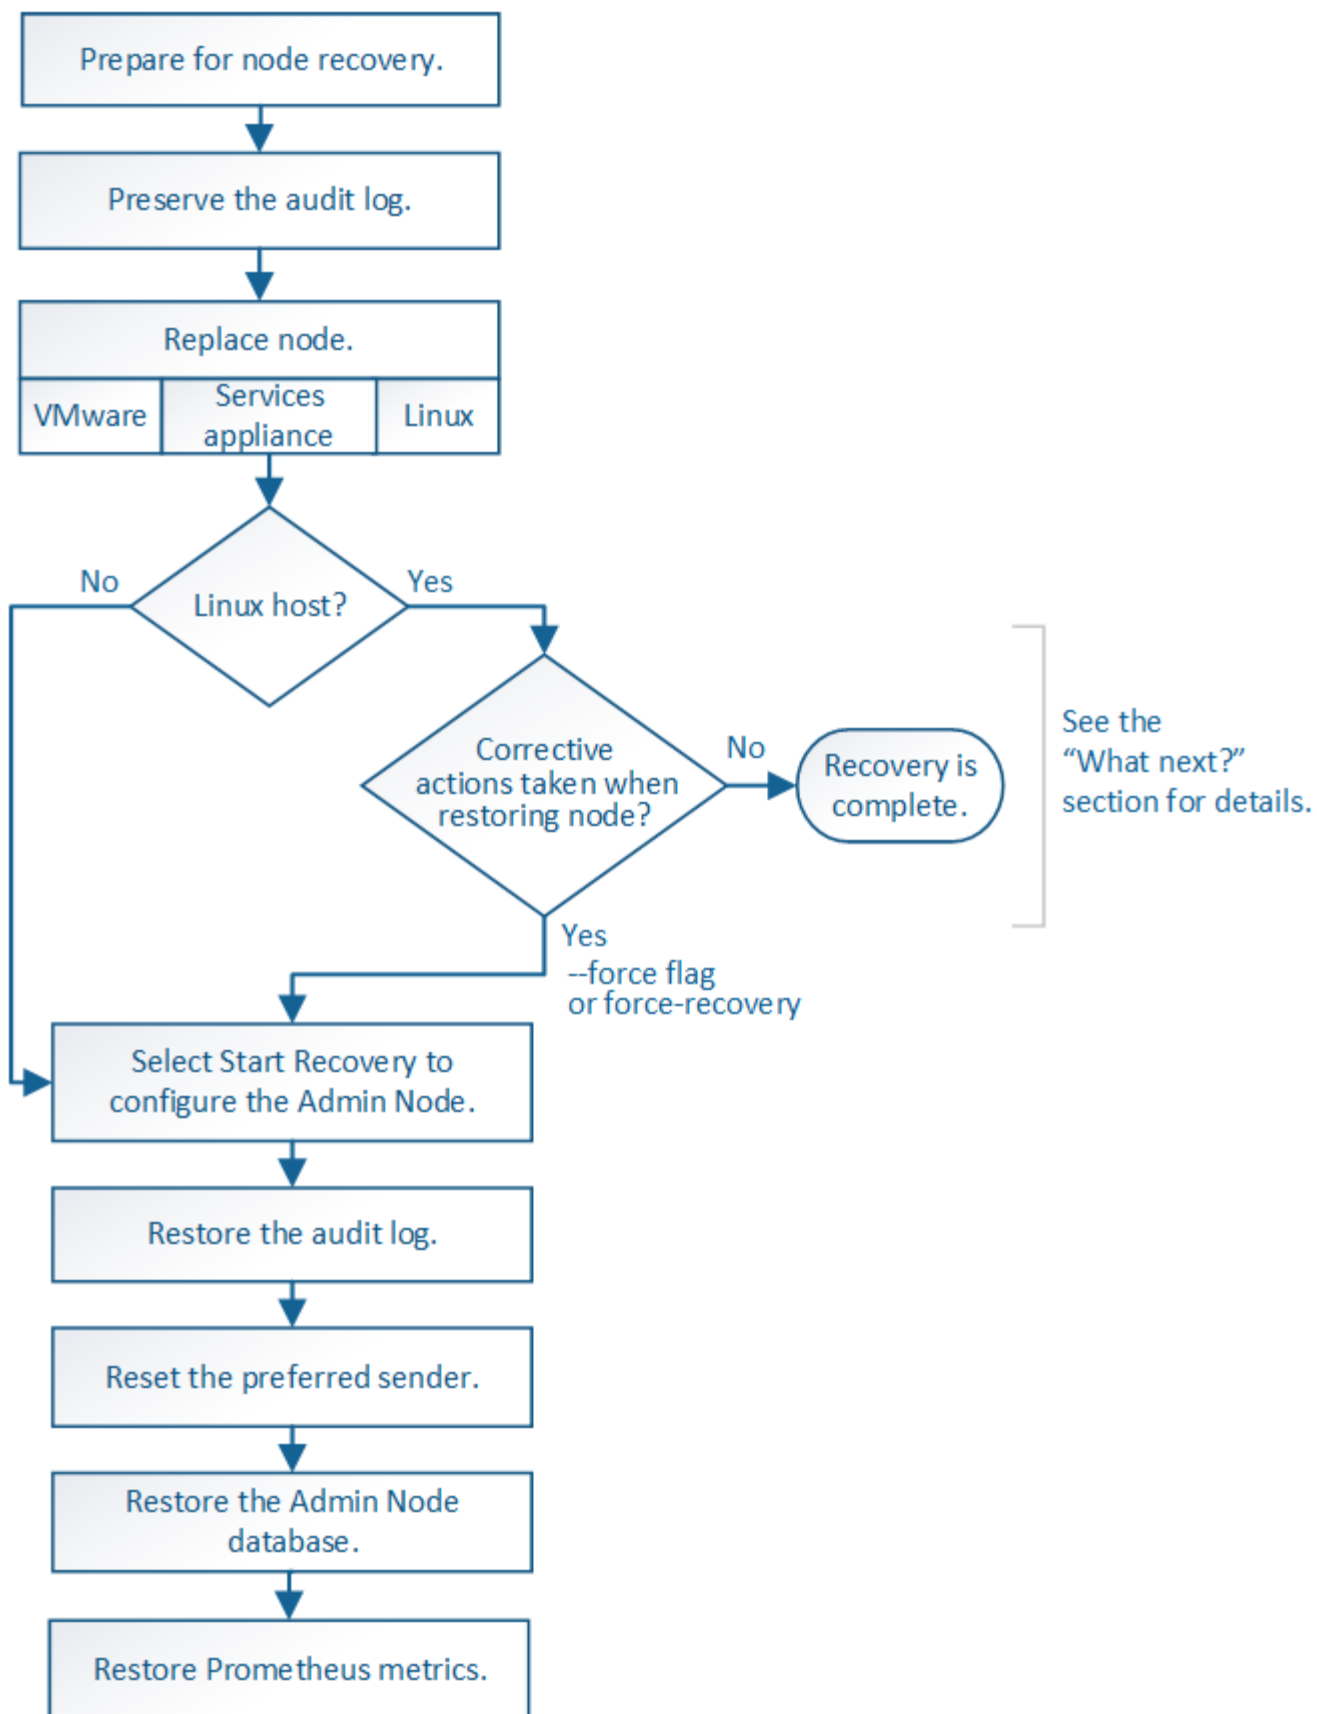
Decommission procedures

You might want to permanently remove grid nodes or an entire data center site from your StorageGRID

system.

For example, you might want to decommission one or more grid nodes in these cases:

- You have added a larger Storage Node to the system and you want to remove one or more smaller Storage Nodes, while at the same time preserving objects.
- You require less total storage.
- You no longer require a Gateway Node or a non-primary Admin Node.
- Your grid includes a disconnected node that you cannot recover or bring back online.

You can use the Decommission Nodes page in the Grid Manager to remove the following types of grid nodes:

- Storage Nodes, unless not enough nodes would remain at the site to support certain requirements
- Gateway Nodes
- Non-primary Admin Nodes

Decommission Nodes

Before decommissioning a grid node, review the health of all nodes. If possible, resolve any issues or alarms before proceeding.

Select the checkbox for each grid node you want to decommission. If decommission is not possible for a node, see the Recovery and Maintenance Guide to learn how to proceed.

Grid Nodes

Search   🔍

	Name		Site		Type		Has ADC	Health	Decommission Possible
	DC1-ADM1		Data Center 1		Admin Node		-	✅	No, primary Admin Node decommissioning is not supported.
	DC1-ARC1		Data Center 1		Archive Node		-	✅	No, Archive Nodes decommissioning is not supported.
☐	DC1-G1		Data Center 1		API Gateway Node		-	✅	✔
	DC1-S1		Data Center 1		Storage Node		Yes	✅	No, site Data Center 1 requires a minimum of 3 Storage Nodes with ADC services.
	DC1-S2		Data Center 1		Storage Node		Yes	✅	No, site Data Center 1 requires a minimum of 3 Storage Nodes with ADC services.
	DC1-S3		Data Center 1		Storage Node		Yes	✅	No, site Data Center 1 requires a minimum of 3 Storage Nodes with ADC services.
☐	DC1-S4		Data Center 1		Storage Node		No	✅	✔
☐	DC2-ADM1		Data Center 2		Admin Node		-	✅	✔
	DC2-S1		Data Center 2		Storage Node		Yes	✅	No, site Data Center 2 requires a minimum of 3 Storage Nodes with ADC services.

You can use the Decommission Site page in the Grid Manager to remove a site. A connected site decommission removes an operational site and preserves data. A disconnected site decommission removes a failed site but does not preserve data. The Decommission Site wizard guides you through the process of selecting the site, viewing site details, revising the ILM policy, removing site references from ILM rules, and resolving any node conflicts.

Network maintenance procedures

Some of the network maintenance procedures you might need to perform include the following:

- Updating the subnets on the Grid Network
- Using the Change IP tool to change the networking configuration that was initially set during grid deployment
- Adding, removing, or updating domain name system (DNS) servers
- Adding, removing, or updating network time protocol (NTP) servers to ensure that data is synchronized accurately between grid nodes
- Restoring network connectivity to nodes that might have become isolated from the rest of the grid

Host-level and middleware procedures

Some maintenance procedures are specific to StorageGRID nodes that are deployed on Linux or VMware, or are specific to other components of the StorageGRID solution. For example, you might want to migrate a grid node to a different Linux host or perform maintenance on an Archive Node that is connected to Tivoli Storage Manager (TSM).

Appliance node cloning

Appliance node cloning lets you easily replace an existing appliance node in your grid with an appliance of newer design or increased capabilities that is part of the same logical StorageGRID site. The process transfers all data to the new appliance, placing it in service to replace the old appliance node and leaving the old appliance in a pre-install state. Cloning provides a hardware-upgrade process that is easy to perform, and provides an alternate method for replacing appliances.

Grid node procedures

You might need to perform certain procedures on a specific grid node. For example, you might need to reboot a grid node or manually stop and restart a specific grid node service. Some grid node procedures can be performed from the Grid Manager; others require you to log in to the grid node and use the node's command line.

Related information

- Administer StorageGRID
- Upgrade software
- Expand your grid
- Recover and maintain

Download the Recovery Package

The Recovery Package is a downloadable .zip file that contains deployment-specific files and software needed to install, expand, upgrade, and maintain a StorageGRID system.

The Recovery Package file also contains system-specific configuration and integration information, including server hostnames and IP addresses, and highly confidential passwords needed during system maintenance, upgrade, and expansion. The Recovery Package is required to recover from the failure of the primary Admin Node.

When installing a StorageGRID system, you are required to download the Recovery Package file and to confirm that you can successfully access the contents of this file. You should also download the file each time the grid topology of the StorageGRID system changes because of maintenance or upgrade procedures.

Recovery Package

Enter your provisioning passphrase and click Start Download to save a copy of the Recovery Package file. Download the file each time the grid topology of the StorageGRID system changes because of maintenance or upgrade procedures, so that you can restore the grid if a failure occurs.

When the download completes, copy the Recovery Package file to two safe, secure, and separate locations.

Important: The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

Provisioning Passphrase   [••••••]

[Start Download]

After downloading the Recovery Package file and confirming you can extract the contents, copy the Recovery Package file to two safe, secure, and separate locations.

> (i) The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

Related information

- Upgrade software
- Expand your grid
- Recover and maintain

Use StorageGRID support options

The Grid Manager provides options to help you work with technical support if an issue arises with your StorageGRID system.

Configure AutoSupport

The AutoSupport feature enables your StorageGRID system to send health and status messages to technical support. Using AutoSupport can significantly speed problem determination and resolution. Technical support can also monitor the storage needs of your system and help you determine if you need to add new nodes or sites. Optionally, you can configure AutoSupport messages to be sent to one additional destination.

You configure AutoSupport using the Grid Manager (**SUPPORT** > **Tools** > **AutoSupport**). The **AutoSupport** page has two tabs: **Settings** and **Results**.

Information included in AutoSupport messages

AutoSupport messages include information such as the following:

- StorageGRID software version
- Operating system version
- System-level and location-level attribute information

- Recent alerts and alarms (legacy system)

- Current status of all grid tasks, including historical data

- Admin Node database usage

- Number of lost or missing objects

- Grid configuration settings

- NMS entities

- Active ILM policy

- Provisioned grid specification file

- Diagnostic metrics

You can enable the AutoSupport feature and the individual AutoSupport options when you first install StorageGRID, or you can enable them later. If AutoSupport is not enabled, a message appears on the Grid Manager Dashboard. The message includes a link to the AutoSupport configuration page.

The AutoSupport feature is disabled. You should enable AutoSupport to allow StorageGRID to send health and status messages to technical support for proactive monitoring and troubleshooting.

If you close the message, it will not appear again until your browser cache is cleared, even if AutoSupport remains disabled.

Use Digital Advisor

Active IQ Digital Advisor is cloud-based and leverages predictive analytics and community wisdom from NetApp's installed base. Its continuous risk assessments, predictive alerts, prescriptive guidance, and automated actions help you prevent problems before they occur, leading to improved system health and higher system availability.

You must enable AutoSupport if you want to use the Digital Advisor dashboards and functionality on the NetApp Support site.

Digital Advisor Documentation

Collect StorageGRID logs

To help troubleshoot a problem, you might need to collect log files and forward them to technical support.

StorageGRID uses log files to capture events, diagnostic messages, and error conditions. The bycast.log file is maintained for every grid node and is the primary troubleshooting file. StorageGRID also creates log files for individual StorageGRID services, log files related to deployment and maintenance activities, and log files related to third-party applications.

Users who have the appropriate permissions and who know the provisioning passphrase for your StorageGRID system can use the Logs page in the Grid Manager to gather log files, system data, and configuration data. When you collect logs, you select a node or nodes and specify a time period. Data is collected and archived in a `.tar.gz` file, which you can download to a local computer. Inside this file, there is one log file archive for each grid node.

Use metrics and run diagnostics

When troubleshooting an issue, you can work with technical support to review detailed metrics and charts for your StorageGRID system. You can also run pre-constructed diagnostic queries to proactively assess key values for your StorageGRID system.

Metrics page

The Metrics page provides access to the Prometheus and Grafana user interfaces. Prometheus is open-source software for collecting metrics. Grafana is open-source software for metrics visualization.

> ⓘ The tools available on the Metrics page are intended for use by technical support. Some features and menu items within these tools are intentionally non-functional and are subject to change.

Metrics

Access charts and metrics to help troubleshoot issues.

> ℹ The tools available on this page are intended for use by technical support. Some features and menu items within these tools are intentionally non-functional.

Prometheus

Prometheus is an open-source toolkit for collecting metrics. The Prometheus interface allows you to query the current values of metrics and to view charts of the values over time.

Access the Prometheus UI using the link below. You must be signed in to the Grid Manager.

- https://   /metrics/graph

Grafana

Grafana is open-source software for metrics visualization. The Grafana interface provides pre-constructed dashboards that contain graphs of important metric values over time.

Access the Grafana dashboards using the links below. You must be signed in to the Grid Manager.

ADE	Grid	S3 - Node
Account Service Overview	ILM	S3 Overview
Alertmanager	Identity Service Overview	S3 Select
Audit Overview	Ingests	Site
Cassandra Cluster Overview	Node	Support
Cassandra Network Overview	Node (Internal Use)	Traces
Cassandra Node Overview	OSL - AsyncIO	Traffic Classification Policy
Cloud Storage Pool Overview	Platform Services Commits	Usage Processing
EC - ADE	Platform Services Overview	Virtual Memory (vmstat)
EC - Chunk Service	Platform Services Processing	
EC Overview	Replicated Read Path Overview	

The link in the Prometheus section of the Metrics page allows you to query the current values of StorageGRID metrics and to view graphs of the values over time.

Metrics that include *private* in their names are intended for internal use only and are subject to change between StorageGRID releases without notice.

The links in the Grafana section of the Metrics page allow you to access pre-constructed dashboards containing graphs of StorageGRID metrics over time.

Diagnostics page

The Diagnostics page performs a set of pre-constructed diagnostic checks on the current state of the grid. In the example, all diagnostics have a Normal status.

Diagnostics

This page performs a set of diagnostic checks on the current state of the grid. A diagnostic check can have one of three statuses:

✔ **Normal**: All values are within the normal range.

⚠ **Attention**: One or more of the values are outside of the normal range.

❌ **Caution**: One or more of the values are significantly outside of the normal range.

Diagnostic statuses are independent of current alerts and might not indicate operational issues with the grid. For example, a diagnostic check might show Caution status even if no alert has been triggered.

[Run Diagnostics]

✔ Cassandra blocked task queue too large	˅
✔ Cassandra commit log latency	˅
✔ Cassandra commit log queue depth	˅
✔ Cassandra compaction queue too large	˅

Clicking a specific diagnostic lets you see details about the diagnostic and its current results.

In this example, the current CPU utilization for every node in a StorageGRID system is shown. All node values are below the Attention and Caution thresholds, so the overall status of the diagnostic is Normal.

✔ CPU utilization

Checks the current CPU utilization on each node.

To view charts of CPU utilization and other per-node metrics, access the **Node** Grafana dashboard.

Status   ✔ Normal

Prometheus query   sum by (instance) (sum by (instance, mode) (irate(node_cpu_seconds_total{mode!="idle"}[5m])) / count by (instance, mode)(node_cpu_seconds_total{mode!="idle"}))

View in Prometheus ⌁

Thresholds   ⚠ Attention >= 75%
❌ Caution   >= 95%

Status	Instance ⇅	CPU Utilization ⇅
✔	DC1-ADM1	2.598%
✔	DC1-ARC1	0.937%
✔	DC1-G1	2.119%
✔	DC1-S1	8.708%
✔	DC1-S2	8.142%
✔	DC1-S3	9.669%
✔	DC2-ADM1	2.515%
✔	DC2-ARC1	1.152%
✔	DC2-S1	8.204%
✔	DC2-S2	5.000%
✔	DC2-S3	10.469%

Related information

- Administer StorageGRID

- Configure network settings