



Using StorageGRID

StorageGRID

NetApp

March 10, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-115/primer/controlling-storagegrid-access.html> on March 10, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Using StorageGRID	1
Controlling StorageGRID access	1
Controlling access to the Grid Manager	1
Enabling single sign-on	1
Using client certificates	2
Changing grid passwords	2
Managing tenants and client connections	2
Tenant accounts	2
Client connections to StorageGRID nodes	3
Configuring network settings	4
Domain names	4
High availability groups	4
Link costs	4
Load balancer endpoints	4
Proxy settings	5
Server certificates	5
Traffic classification policies	5
Untrusted Client Networks	5
Configuring system settings	6
Display options	6
Grid options	6
Key management servers	6
S3 Object Lock	6
Storage options	7
Using information lifecycle management	7
Example ILM rule	8
How an ILM policy evaluates objects	8
Example ILM policy	9
Monitoring StorageGRID operations	10
Viewing the Nodes page	10
Monitoring and managing alerts	15
Using SNMP monitoring	17
Reviewing audit messages	17
Performing maintenance procedures	18
Software updates	18
StorageGRID software upgrades	18
Hotfixes	18
SANtricity OS upgrades	19
Expansion procedures	20
Node recovery procedures	21
Decommission procedures	22
Network maintenance procedures	24
Host-level and middleware procedures	24

Appliance node cloning	24
Grid node procedures	25
Downloading the Recovery Package	25
Using StorageGRID support options	26
Configuring AutoSupport	26
Collecting StorageGRID logs	28
Using metrics and running diagnostics	29

Using StorageGRID

After you install grid nodes and StorageGRID networks, you can begin to configure and use StorageGRID. Some of the tasks you will perform include controlling user access to system administration functions, setting up tenant accounts , managing client connections, setting configuration options, managing object locations with ILM, monitoring the health and day-to-day activities of your StorageGRID system, and performing routine and non-routine maintenance activities.

- [Controlling StorageGRID access](#)
- [Managing tenants and client connections](#)
- [Configuring network settings](#)
- [Configuring system settings](#)
- [Using information lifecycle management](#)
- [Monitoring StorageGRID operations](#)
- [Performing maintenance procedures](#)
- [Using StorageGRID support options](#)

Controlling 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.

Controlling 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.



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.

Enabling 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.

Using client certificates

You can use client certificates to allow authorized external clients to access the StorageGRID Prometheus database. Client certificates provide a secure way to use external tools to monitor StorageGRID. You can provide your own client certificate or generate one using the Grid Manager.

Changing 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](#)

Managing 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. 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).
2. 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.
3. 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.

Related information

[Administer StorageGRID](#)

[Use a tenant account](#)

[Use S3](#)

[Use Swift](#)

Configuring 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

High availability groups use virtual IP addresses (VIPs) to provide active-backup access to Gateway Node or Admin Node services. An HA group consists of one or more network interfaces on Admin Nodes and Gateway Nodes. When creating an HA group, you select network interfaces belonging to the Grid Network (`eth0`) or the Client Network (`eth2`).



The Admin Network does not support HA VIPs.

An HA group maintains one or more virtual IP addresses that are added to the active interface in the group. If the active interface becomes unavailable, the virtual IP addresses are moved to another interface. This failover process generally takes only a few seconds and is fast enough that client applications should experience little impact and can rely on normal retry behaviors to continue operation.

You might want to use high availability (HA) groups for several reasons.

- An HA group can provide highly available administrative connections to the Grid Manager or the Tenant Manager.
- An HA group can provide highly available data connections for S3 and Swift clients.
- An HA group that contains only one interface allows you to provide many VIP addresses and to explicitly set IPv6 addresses.

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.

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.



Server certificates

You can upload two types of server certificates:

- Management Interface Server Certificate, which is the certificate used for accessing the management interface.
- Object Storage API Service Endpoints Server Certificate, which secures the S3 and Swift endpoints for connections directly to Storage Nodes or when using the CLB service on a Gateway Node.



The CLB service is deprecated.

Load balancer certificates are configured on the Load Balancer Endpoints page. Key management server (KMS) certificates are configured on the Key Management Server page.

Traffic classification policies

Traffic classification policies allow you to create rules for identifying and handling 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.

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](#)

[Managing tenants and client connections](#)

Configuring 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.

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.

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.



Storage options

Storage options allow you to control object segmentation and to define storage watermarks to manage a Storage Node's usable storage space.

Using 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.

Two copies at two sites for Tenant A

Description:

Applies only to Tenant A

Ingest Behavior:

Balanced

Tenant Accounts:

Tenant A (34176783492629515782)

Reference Time:

Ingest Time

Filtering Criteria:

Matches all objects.

Retention Diagram:

Trigger

Day 0

Site 1

Site 2

Duration

Forever

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.

Example ILM policy

This example ILM policy uses three ILM rules.

Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Example ILM policy

Reason for change

New policy

Rules

1. Select the rules you want to add to the policy.
2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules				
	Default	Rule Name	Tenant Account	Actions
⬇		Rule 1: 3 replicated copies for Tenant A	Tenant A (58889986524346589742)	✕
⬇		Rule 2: Erasure coding for objects greater than 1 MB	—	✕
	✓	Rule 3: 2 copies 2 data centers (default)	—	✕


Cancel

Save

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 larger 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](#)


Monitoring StorageGRID operations

The Grid Manager provides information for monitoring the daily activities of your StorageGRID system, including its health.

- [Viewing the Nodes page](#)
- [Monitoring and managing alerts](#)
- [Using SNMP monitoring](#)
- [Reviewing audit messages](#)

Viewing 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.



From the tree view on the left, you can see all the sites and all the nodes in your StorageGRID system. The icon for each node indicates if the node is connected or if there are any active alerts.


Connection state icons

If a node is disconnected from the grid, the tree view shows a blue or gray connection state icon, not the icon for any underlying alerts.

- **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.

Alert icons

If a node is connected to the grid, the tree view shows one of the following icons, depending on if there are any current alerts for the node.

- **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.
- **Normal** ✅: No alerts are active, and the node is connected to the grid.

Viewing details for a system, site, or node

To view the available information, click the appropriate links on the left, 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.

Related information

[Monitor & troubleshoot](#)

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	<ul style="list-style-type: none">• Provides basic information about each node.• Shows any current, unacknowledged alarms affecting the node.	All nodes
Hardware	<ul style="list-style-type: none">• 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.	All nodes, each site, and the entire grid
Storage	<ul style="list-style-type: none">• 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
Events	Displays a count of any system error or fault event, including errors such as network errors.	All nodes


Tab name	Description	Included for
Objects	<ul style="list-style-type: none"> 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
ILM	<p>Provides information about Information Lifecycle Management (ILM) operations.</p> <ul style="list-style-type: none"> 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	<p>Includes performance and diagnostic graphs related to the Load Balancer service.</p> <ul style="list-style-type: none"> 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 less than 8.70.

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 (**Help > 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.



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.



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

Related information

[Monitoring and managing alerts](#)

[Using StorageGRID support options](#)

[Monitor & troubleshoot](#)

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 & troubleshoot](#)

Monitoring and managing 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.

Current Alerts [Learn more](#)

View the current alerts affecting your StorageGRID system.

							<input checked="" type="checkbox"/> Group alerts <div>Active</div>
Name	Severity	Time triggered	Site / Node	Status	Current values		
Unable to communicate with node One or more services are unresponsive or cannot be reached by the metrics collection job.	2 Major	9 minutes ago <i>(newest)</i> 19 minutes ago <i>(oldest)</i>		2 Active			
Low root disk capacity The space available on the root disk is low.	Minor	25 minutes ago	Data Center 1 / DC1-S1-99-51	Active	Disk space available: 2.00 GB Total disk space: 21.00 GB		
Expiration of server certificate for Storage API Endpoints The server certificate used for the storage API endpoints is about to expire.	Major	31 minutes ago	Data Center 1 / DC1-ADM1-99-49	Active	Days remaining: 14		
Expiration of server certificate for Management Interface The server certificate used for the management interface is about to expire.	Minor	31 minutes ago	Data Center 1 / DC1-ADM1-99-49	Active	Days remaining: 30		
Low installed node memory The amount of installed memory on a node is low.	8 Critical	a day ago <i>(newest)</i> a day ago <i>(oldest)</i>		8 Active			

- 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



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

Managing 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

[Monitor & troubleshoot](#)

Using 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. Informes 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.



SNMP notifications are not sent for every alarm or every alarm severity.

Related information

[Monitor & troubleshoot](#)

Reviewing 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.

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

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 the instructions for administering StorageGRID.

Related information

[Review audit logs](#)

[Administer StorageGRID](#)

Performing maintenance procedures

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.

StorageGRID Hotfix

Before starting the hotfix process, you must confirm that there are no active alerts and that all grid nodes are online and available.

When the primary Admin Node is updated, services are stopped and restarted. Connectivity might be interrupted until the services are back online.

Hotfix file

Hotfix file ?

Browse

Passphrase

Provisioning Passphrase ?

Start

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

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

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 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

i 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 NodesIn 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	<div><div></div></div>	Waiting for NTP to synchronize
DC2-S1-185	Site A	172.17.3.185/21	<div><div></div></div>	Waiting for Dynamic IP Service peers
DC2-S2-186	Site A	172.17.3.186/21	<div><div></div></div>	Waiting for NTP to synchronize
DC2-S3-187	Site A	172.17.3.187/21	<div><div></div></div>	Waiting for NTP to synchronize
DC2-S4-188	Site A	172.17.3.188/21	<div><div></div></div>	Waiting for Dynamic IP Service peers
DC2-ARC1-189	Site A	172.17.3.189/21	<div><div></div></div>	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:

















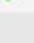
- Storage Nodes, unless not enough nodes would remain at the site to support certain requirements
- Gateway Nodes
- Non-primary Admin Nodes

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.	
<input type="checkbox"/>	DC1-ADM2	Data Center 1	Admin Node	-			
<input type="checkbox"/>	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.	
<input type="checkbox"/>	DC1-S4	Data Center 1	Storage Node	No			
<input type="checkbox"/>	DC1-S5	Data Center 1	Storage Node	No			

Passphrase

Provisioning
Passphrase

Start Decommission

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.

Decommission Site



When you decommission a site, all nodes at the site and the site itself are permanently removed from the StorageGRID system.

Review the table for the site you want to remove. If Decommission Possible is Yes, select the site. Then, select **Next** to ensure that the site is not referred to by ILM and that all StorageGRID nodes are in the correct state.

You might not be able to remove certain sites. For example, you cannot decommission the site that contains the primary Admin Node or a site that contains an Archive Node.

Sites

	Site Name	Used Storage Capacity ?	Decommission Possible
<input checked="" type="radio"/>	Raleigh	3.93 MB	✓
<input checked="" type="radio"/>	Sunnyvale	3.97 MB	✓
<input type="radio"/>	Vancouver	3.90 MB	No. This site contains the primary Admin Node.

Next

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 (source) in your grid with a compatible appliance (target) 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](#)

[Maintain & recover](#)

Downloading 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.



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](#)

Using StorageGRID support options

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

Configuring 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.

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
- Events information as listed on the **Nodes** > **node** > **Events** page
- 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 ManagerDashboard. 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.



You can select the “x” symbol  to close the message. The message will not appear again until your browser cache is cleared, even if AutoSupport remains disabled.

Using Active IQ

Active IQ is a cloud-based digital advisor that 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 Active IQ dashboards and functionality on the NetApp Support site.

[Active IQ Digital Advisor Documentation](#)

Accessing AutoSupport settings

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

AutoSupport

The AutoSupport feature enables your StorageGRID system to send periodic and event-driven health and status messages to technical support to allow proactive monitoring and troubleshooting. StorageGRID AutoSupport also enables the use of Active IQ for predictive recommendations.

Settings Results

Protocol Details

Protocol ? ☒ HTTPS ☐ HTTP ☐ SMTP

NetApp Support Certificate Validation ? Use NetApp support certificate ▼

AutoSupport Details

Enable Weekly AutoSupport ? ☒

Enable Event-Triggered AutoSupport ? ☒

Enable AutoSupport on Demand ? ☐

Additional AutoSupport Destination

Enable Additional AutoSupport Destination ? ☐

Save Send User-Triggered AutoSupport

Protocols for sending AutoSupport messages

You can choose one of three protocols for sending AutoSupport messages:

- HTTPS
- HTTP
- SMTP

If you send AutoSupport messages using HTTPS or HTTP, you can configure a non-transparent proxy server between Admin Nodes and technical support.

If you use SMTP as the protocol for AutoSupport messages, you must configure an SMTP mail server.

AutoSupport options

You can use any combination of the following options to send AutoSupport messages to technical support:

- **Weekly:** Automatically send AutoSupport messages once per week. Default setting: Enabled.
- **Event-triggered:** Automatically send AutoSupport messages every hour or when significant system events occur. Default setting: Enabled.
- **On Demand:** Allow technical support to request that your StorageGRID system send AutoSupport messages automatically, which is useful when they are actively working an issue (requires HTTPS AutoSupport transmission protocol). Default setting: Disabled.
- **User-triggered:** Manually send AutoSupport messages at any time.

Related information

[Administer StorageGRID](#)

[Configuring network settings](#)

Collecting StorageGRID logs

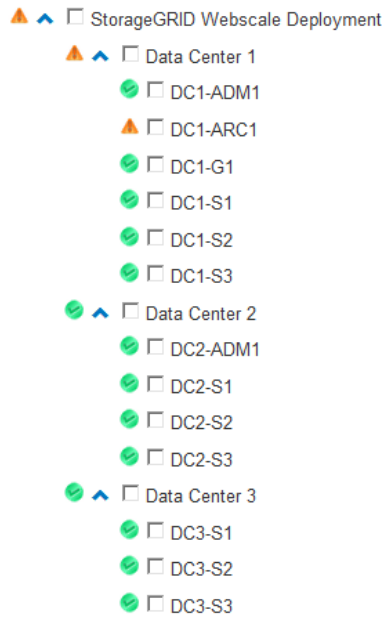
To help troubleshoot a problem, you might need to collect log files and forward them to technical support.

















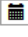
StorageGRID uses log files to capture events, diagnostic messages, and error conditions. The bycast.log file is maintained for every grid node and is the primary troubleshooting file. StorageGRID also creates log files for individual StorageGRID services, log files related to deployment and maintenance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

Logs


Collect log files from selected grid nodes for the given time range. Download the archive package after all logs are ready.



-  ☐ StorageGRID Webscale Deployment
 -  ☐ Data Center 1
 -  ☐ DC1-ADM1
 -  ☐ DC1-ARC1
 -  ☐ DC1-G1
 -  ☐ DC1-S1
 -  ☐ DC1-S2
 -  ☐ DC1-S3
 -  ☐ Data Center 2
 -  ☐ DC2-ADM1
 -  ☐ DC2-S1
 -  ☐ DC2-S2
 -  ☐ DC2-S3
 -  ☐ Data Center 3
 -  ☐ DC3-S1
 -  ☐ DC3-S2
 -  ☐ DC3-S3

Log Start Time

2018-04-18



01

:


38

PM

MDT

Log End Time

2018-04-18



05

:

38

PM

MDT

Notes

Provisioning
Passphrase

Collect Logs

Related information

[Monitor & troubleshoot](#)

[Administer StorageGRID](#)

Using metrics and running 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://storagegrid.com/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	Node
Account Service Overview	Node (Internal Use)
Alertmanager	Platform Services Commits
Audit Overview	Platform Services Overview
Cassandra Cluster Overview	Platform Services Processing
Cassandra Network Overview	Replicated Read Path Overview
Cassandra Node Overview	S3 - Node
Cloud Storage Pool Overview	S3 Overview
EC - ADE	Site
EC - Chunk Service	Support
Grid	Traces
ILM	Traffic Classification Policy
Identity Service Overview	Usage Processing
Ingests	Virtual Memory (vmstat)

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.

PrometheusAlertsGraphStatus ▾Help

☐ Enable query history

Expression (press Shift+Enter for newlines)

Execute

- insert metric at cursor - ▾

GraphConsole

Element	Value
no data	

Remove Graph

Add Graph



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

[Monitor & troubleshoot](#)

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.